From Plausible Tomorrows to Prompt Action on AI

Dealing with AI-Augmented Risks to Dutch National Security

Alexandre Gomes
Koen Aartsma
Maaike Okano-Heijmans
Jelle van den Wijngaard

Clingendael Report

Clingendael
Netherlands Institute of International Relations

Clingendael

Netherlands Institute of International Relations

From Plausible Tomorrows to Prompt Action on AI

Dealing with AI-Augmented Risks to Dutch National Security

Alexandre Gomes
Koen Aartsma
Maaike Okano-Heijmans
Jelle van den Wijngaard

Clingendael Report
November 2025

Disclaimer and Acknowledgements
Research for, and the production of, this Clingendael Report were conducted for the General Intelligence and Security Service (AIVD) of the Netherlands and the Dutch Authority for Digital Infrastructure (RDI). The authors sincerely thank the members of AI36 and of Clingendael's General Purpose Artificial Intelligence Council (GPAI Raad) for their valuable input and comments on an earlier draft of this report. Responsibility for the content and the opinions expressed rests solely with the authors.

November 2025

About the Clingendael Institute

The Netherlands Institute of International Relations 'Clingendael' is a leading think tank and academy on international affairs. Through our analyses, training and public platform activities we aim to inspire and equip governments, businesses, and civil society to contribute to a secure, sustainable and just world.

The Clingendael Institute
P.O. Box 93080
2509 AB The Hague
The Netherlands

Follow us on social media
@clingendaelorg
The Clingendael Institute
The Clingendael Institute
clingendael_institute
Clingendael Institute

Email: [info@clingendael.org / cru@clingendael.org]
Website: [www.clingendael.org / www.clingendael.org/cru]

About the authors

Alexandre Gomes is a Research Fellow at the Netherlands Institute of International Relations 'Clingendael' in The Hague, where he is part of the EU and Global Affairs Unit and of the 'Geopolitics of Technology and Digitalisation' programme.

Koen Aartsma is Head of the Strategic Foresight and Intelligence Programme and a Senior Research Fellow at Clingendael's Security Unit. He leads Clingendael's geopolitical trend analysis, scenario development and horizon scanning. This includes Clingendael's projects and strategic foresights for the Netherlands' Ministries of Foreign Affairs and Defence, as well as foresights for the Ministry of Justice and Security.

Maaike Okano-Heijmans is a Senior Research Fellow at the Netherlands Institute of International Relations 'Clingendael' in The Hague, where she leads the 'Geopolitics of Technology and Digitalisation' programme. She is also a Visiting Lecturer in the Master of Science in International Relations and Diplomacy (MIRD) programme at the University of Leiden.

Jelle van den Wijngaard is a Research Fellow at the Netherlands Institute of International Relations 'Clingendael' in The Hague, where he is part of the EU and Global Affairs Unit and of the 'Geopolitics of Technology and Digitalisation' programme.

Contents

Executive Summary

Artificial Intelligence (AI) is redefining technological, economic and security landscapes. The transformative technology introduces new levels of disruption and reorders traditional power balances. At the same time, it offers possibilities for increased productivity, efficiency and automation. The Netherlands and the European Union (EU) face the dual challenge of leveraging AI's potential while safeguarding national security, democratic resilience and strategic autonomy in an environment mostly dominated by the United States and China.

Building on ten foundational insights from expert conversations – capturing the political, technical and societal dynamics shaping today's AI landscape (the 'baseline') – this report explores three scenarios for how AI could affect Dutch national security over the next five years. These scenarios – so-called 'Plausible Tomorrows' – rest on two key uncertainties: (1) the extent to which the Netherlands, the EU and the public can access advanced AI, for example via open-source models and proliferation; and (2) how ubiquitous the technology becomes, from its integration into everyday devices to whether agentic AI[1] fulfils its promises.

Plausible Tomorrow 1 envisions widespread proliferation of open-source AI, empowering both state and non-state actors, but fuelling disinformation, cyberattacks and semi-autonomous terrorism. Plausible Tomorrow 2 highlights the rise of agentic AI and its disruptive integration into all spheres, intensifying geopolitical competition, labour market upheaval and social unrest. Plausible Tomorrow 3 considers a bifurcated world where the EU's digital dependency deepens and the bloc comes to rely fully on either US or Chinese AI systems, effectively forfeiting its digital sovereignty and resilience irreversibly.

Across all scenarios, AI destabilises the information environment and erodes trust in democratic institutions; creates risks and accelerates vulnerabilities in critical sectors such as defence and health; and imposes significant ecological and energy burdens. At the same time, AI offers opportunities in sectors where

1 Agentic AI is an artificial intelligence system that can act autonomously towards specific goals, making decisions or taking actions without direct human commands.

the Netherlands holds distinct strengths, or niches, such as AgriTech and semiconductors.

Meeting this challenge requires strategic foresight rooted in, and guiding, a long-term strategic vision for AI; proactive policymaking that sets the conditions for European companies to thrive; and investment at a scale that brings Dutch and European research and development (R&D) spending in line with leading peers, namely at or above the 3 per cent of GDP benchmark – and all while strengthening national security and resilience.

Key recommendations that emerge from our analysis of the three Plausible Tomorrows include:

- Control dual-use autonomous systems to mitigate security and ethical risks;
- Mandate secure and reliable update practices across AI systems;
- Safeguard the information environment against manipulation and misuse;
- Prioritise niche domains where the Netherlands already holds a competitive edge;
- Protect human-critical, non-negotiable domains against harmful algorithmic outcomes;
- Strengthen the competitiveness and indispensability of strategic industries in the Netherlands and EU through domestic and sovereign computational power, resilient deep-tech supply chains and EU-aligned funding.

In short, safeguarding democratic institutions, strengthening Europe's AI capabilities and shaping global standards for responsible AI are essential for the Netherlands and the EU to secure their place in an AI-driven future.

Introduction

Artificial Intelligence (AI) is rapidly transforming global technological landscapes. The transformative technology introduces new levels of disruption and reorders traditional power balances. At the same time, it offers possibilities for increased productivity, efficiency and automation. AI's far-reaching implications span the societal, economic, political and military domains, with direct consequences for national security.

This report assesses how AI developments may threaten national security interests, while recognising that responses often require both Dutch and EU-level action. In line with the 2023 Dutch National Security Strategy, we understand national security as the protection of the Netherlands' six national security interests:[2] safeguarding its territory (territorial security); protecting its people (physical safety); the functioning of critical societal and economic sectors (economic security); the resilience of its living environment (ecological security); the democratic rule of law and social cohesion (social and political stability); and the country's capacity to act autonomously in international affairs and uphold the international legal order and stability.

The current AI boom and increasing availability of AI software and open-source models is facilitating the proliferation and potential democratisation of AI technologies. The AI models that are most used by consumers – especially Large Language Models (LLMs) – are 'proprietary models', as they belong to a few (mostly) American companies. A shift towards open-source models – of varying quality, safety and bias – could help reduce Europe's digital dependencies and thereby contribute to addressing national security concerns. After all, both state and non-state actors increasingly have access to powerful tools capable of shaping public discourse, executing cyber operations and deploying autonomous systems.

2 Government of the Netherlands, *Security Strategy for the Kingdom of the Netherlands*, 3 April 2023. Please note that our analysis of each of the six national security interests is not exhaustive – that is, not all themes are addressed in all the analysis sections of this report.

In this volatile context, foresight exercises, including scenario building, can aid in imagining and preparing for a range of so-called 'Plausible Tomorrows', each with a unique set of challenges and opportunities. Against this backdrop, this report asks: How might AI developments threaten or strengthen the Netherlands' national security interests, and what actions can Dutch and EU policymakers take to mitigate risks and seize opportunities? To address this question, we employ a foresight approach that uses scenario building to explore 'plausible tomorrows of AI's impact on national security.[3]

The report proceeds as follows: it first introduces the analytical framework and key trends shaping AI developments; then presents three Plausible Tomorrows that explore possible futures; and finally distils policy implications and recommendations for Dutch and EU decision-makers.

On the Global Chessboard: Competing for AI Power

While Generative AI, for instance LLMs, has captured public attention since the launch of the ChatGPT chatbot at the end of 2022, AI's disruptive potential goes well beyond natural language processing. Progress in computer vision and autonomous robotics, for instance, supports new capabilities in surveillance, battlefield automation and decision-making systems. The impact is currently most visible in battlefield operations, in theatres like Ukraine or the Middle East, but also in the United States (US), where big-data analytics and integration company Palantir promises to take surveillance a step further by integrating data from multiple federal agencies,[4] and in China, which is emerging in pole position in the development and commercialisation of humanoid robots.[5]

Amid these trends, leading powers the United States and China have adopted proactive strategies to secure their positions. The Trump administration's AI Action Plan of July 2025 adopts a largely hands-off, pro-innovation approach to governing AI and calls for the US to 'win the race' in this rapidly evolving

3 This report relies on input from a scenario workshop organised in May 2025 to gain insights from experts in AI.

4 *New York Times*, 'Trump Taps Palantir to Compile Data on Americans', 30 May 2025.

5 *South China Morning Post*, 'World Robotics Conference in China Marks 10th Year with JD.com as Strategic Partner', 7 August 2025.

technology and related industries.[6] The US and China lead on frontier models (mostly LLMs) and applications (with the US still leading) as well as in AI-related research and development and innovation (R&D&I, where China has taken the lead).[7] New players in the Middle East and the Gulf, such as Israel, Saudi Arabia and the United Arab Emirates, are entering agreements with companies like Google and chip design giant Nvidia to get a foot in the door.[8]

The European Union (EU) and leading European countries are dealing with limited access to capital and computing power, as well as regulatory dilemmas and industrial gaps. As frontrunners, the United Kingdom and France do hold key AI software (such as British Synthesia) and a leading LLM foundational model (French Mistral's Le Chat). The Netherlands also holds relative AI software strengths, including on the infrastructure side (Axelera AI and Nebius AI) and in sectors such as AgriTech and HealthTech (for example, Nedap and Philips). Even when taken together, however, European countries lag behind the US and China.[9]

Set against this context, the EU positioned itself as the regulatory standard-setter for AI: the world's first binding regulation for ethical AI, the AI Act, has been in effect since August 2025. More recently, the EU and its member states have also been aiming to enhance the bloc's competitiveness in the field. Recent and upcoming EU initiatives like InvestAI, the AI Continent Invest Plan and the Cloud and AI Developments Act include investments in computational power (hereafter 'compute'), access to finance for European start-up and scale-up companies, and commercialisation of R&D&I. The success of such initiatives and the ingenuity of European citizens and businesses will be indispensable in a future where AI systems are expected to play a central role in military deterrence, economic competitiveness and potentially also for work automation in aging societies.

6 The White House, 'Winning the Race: America's AI Action Plan', 23 July 2025.

7 *Business Insider*, 'There's a Key Difference in How China and the US are Integrating their Latest AI Models into Consumer Tech', 11 April 2025.

8 Atlantic Council, 'Advancing US National Security through Middle East AI Negotiations', 9 December 2024. See also, among others, *BBC*, 'Tech Giants are Putting $500bn into "Stargate" to Build Up AI in US', 22 January 2025.

9 Mario Draghi, *The Future of European Competitiveness*, 2024.

Plausible Tomorrows Guiding Action Today

The nature of AI-induced disruptions, as well as asymmetric access to advanced AI technologies, create real and present challenges for the Dutch national security apparatus, today and tomorrow. Starting with a baseline assessment of present-day AI technologies and applications based on expert interviews, this report will introduce three scenarios or 'Plausible Tomorrows', intended to help with imagining potential AI-induced national security risks. An analysis of the national security implications for the Netherlands and the EU follows the description of each plausible tomorrow.

Taking Stock: Where Does AI Stand?

Since 2023, there has been remarkable growth in AI capabilities, which became widely available to the public in the progress made by tools like OpenAI's ChatGPT, Microsoft's Copilot or Google's Gemini. As well as generating text, LLMs are also capable of using other generative AI models, such as for image or video generation. Moreover, tools such as Perplexity incorporate internet connectivity for real-time web search and provide accurate sources with their responses. Besides the deployment of more powerful LLMs, AI systems are sometimes allowed to make decisions on their own, granting them a degree of autonomy or agency. Such AI systems are aptly named 'agentic AI systems'.

The combination of autonomous decision-making and the accelerating pace of innovation creates opportunities for transformative applications, including AI-driven energy-grid management to balance renewable supply and demand, and scientific discovery engines capable of autonomously generating and testing hypotheses. But they also introduce new challenges, for instance in assigning liability in court or challenging unfair or biased decisions. The rapid AI developments affect private companies and public functions in critical domains such as cybersecurity, defence and public infrastructure.

Crucially, all recent AI breakthroughs have been in the domain of so-called 'narrow AI.' That is, they are essentially sophisticated statistical models. They are increasingly good at predicting outcomes in a certain narrowly specified domain, for which they have been 'trained' by inferring logic and rules from large datasets. There are heavy debates on the future of AI and whether 'Artificial General Intelligence' – loosely defined as an AI that is at least as intelligent and adaptable as humans across most tasks – will emerge soon. However, the potential realisation of such technologies is explicitly outside the scope of this paper, which discusses plausible tomorrows extrapolating from current trends in narrow AI developments.

To ensure that the Netherlands remains secure and sovereign in this evolving context, it is vital, first, to take stock of where AI developments stand – from political, technical and societal points of view – to establish a baseline.[10]

Baseline and Trajectory of AI – Ten Key Takeaways

To understand the strategic implications of AI for the EU and the Netherlands, we outline ten foundational insights gained from our expert conversations that capture the political, technical and societal dynamics shaping the current AI landscape.[11]

1. Great powers – notably the US and China, and to a lesser extent the EU and some Member States, particularly France – regard **AI as a strategic race**. Access to infrastructure, talent and funding are considered core requirements for AI leadership.
2. **Uncertainty in time horizons increases.** While AI timelines seem to accelerate, it is unclear where we are on the so-called S-curve of innovation – that is, how mature AI technologies are and where there is still space to improve. This makes prediction impossible and scenario planning very difficult, but indispensable.
3. The **accessibility paradigm is shifting** – on the one hand, new entrants like xAI have rapidly approached top-tier capabilities based on an intense capital investment; on the other hand, DeepSeek has shown that capabilities are not a linear function on compute, and that efficiency using resources is also an element to consider.
4. **The real risk is not AI itself, but how governments, companies and citizens embed AI into society**.[12] Its use can lead to significant efficiency gains, but also amplifies important existing problems, such as the spread of mis- and disinformation, inequality and social instability. For instance, traditional media that had already been under pressure from a changing media landscape are losing additional income as Google's AI-generated search

10 This baseline is drawn from expert interviews, literature analysis and foresight exercises conducted during a Clingendael scenario workshop on 8 April 2025.

11 Please note that the items on this list are not ranked by importance.

12 See *Forbes*, 'MIT Finds 95% of GenAI Pilots Fail Because Companies Avoid Friction', 26 August 2025.

results limit redirects to news pages, which generated an income through advertisements.[13]

5. Proliferation of AI inevitably comes at a security cost. Wider availability of powerful AI through open models **increases the risk of misuse by malign actors**, including in the cyber and biotech domains. The importance of AI safety and oversight of those trends will only increase.
6. The rise of **agentic AI changes the scale and autonomy of AI** use. These systems can plan and act independently, increasing unpredictability.
7. **There is AI beyond LLMs**. Future breakthroughs from AI will also come from less-discussed model types (for example, computer vision, human biology models, semantic and logic models), increasing the complexity and uncertainty in AI.
8. **Questions around model dependencies and trust arise** as winners and losers emerge. The use of foreign models, or models from adversaries, raises concerns over software supply-chain integrity, hidden vulnerabilities and embedded biases.
9. As AI improves, **compute-based risk classification**,[14] such as for General Purpose Artificial Intelligence (GPAI) in the EU's AI Act, will be rendered useless within a short time span. Regulatory benchmarks based on computing power are quickly outdated because of fast-evolving optimisation methods. In other words, equally powerful models can emerge using a decreasing computational volume.
10. **Ecological and energy burdens** are a structural constraint. Training and operating advanced AI consumes vast amounts of electricity, water and rare materials. These pressures are not only a side-effect of proliferation but an inherent feature of scale, creating new dependencies on energy systems and critical resources.[15]

---

13 Compounding the problem, the AI systems that Google and others deploy have been trained on data sourced from the same media outlets now under threat. This essentially happens without the media outlets' consent, which has led, for instance, the *New York Times* to sue OpenAI and Microsoft. This development presents a serious threat for access to information and, ultimately, democracy itself.

14 Compute-based risk classification refers to the practice of assessing AI risk based on the amount of computational power (e.g. processing capacity or training compute) used to develop or operate a model.

15 Several technology firms have signed long-term deals to use nuclear energy to supply power to data-centre operations. While many projects are not yet fully operational, this trend suggests that corporations are exploring alternative energy-supply sources, with potential geopolitical and power balance implications. See, for example, *Reuters*, 'Google Announces Tennessee as Site for Small Modular Nuclear Reactor', 19 August 2025.

Key Uncertainties: AI Access and Ubiquitousness

The baseline above shows the multi-pronged nature of the challenge for the EU and its Member States in relation to AI. Against this backdrop, this report assesses the potential impact of AI on national security based on two main uncertainties: (1) the degree to which the Netherlands, the EU and the public have access to cutting-edge AI capabilities, for example via open source AI proliferation or innovation; and (2) the ubiquitousness of the technology, for example depending on its incorporation into devices and the extent to which agentic AI lives up to the promises made by tech companies. These uncertainties arise as key determinants because they directly shape the power dynamics between state and non-state actors and the extent to which the Netherlands can act independently or must rely on external actors for AI capabilities.

In short, these are key 'known unknowns' with which the EU and its Member States must reckon. While these uncertainties make it difficult to plan for a single future, strategic foresight can help in preparing for a range of plausible developments and outcomes. As such, this report relies on input from a scenario workshop organised in May 2025 to gain insights from experts in AI. For more information on this research method, see the appendix.

The next section will dive into three Plausible Tomorrows, over a three- to five-year horizon. These Plausible Tomorrows, while speculative, engage in different ways with the two key uncertainties highlighted above: access to the latest AI technology and the extent to which AI becomes ubiquitous.[16] Figure 1 shows how the three Plausible Tomorrows interact with these uncertainties. An analysis of the risks and threats follows each scenario.

16 Of course, many other futures can be envisioned, for example: an AI bubble bursting, geopolitical dominance by either China or the United States, breakthroughs in quantum computing that relegate the current AI paradigm to a secondary role, or shifts along the open–closed model spectrum such as open weights. Moreover, a detailed discussion on aspects such as the importance of data ownership, interoperability and AI governance lie beyond the scope of this report because of space constraints, but they merit further analysis elsewhere.

Figure 1 Graphical representation of how the three Plausible Tomorrows interact with the two key uncertainties identified: access to the latest AI technology and the extent to which AI becomes ubiquitous

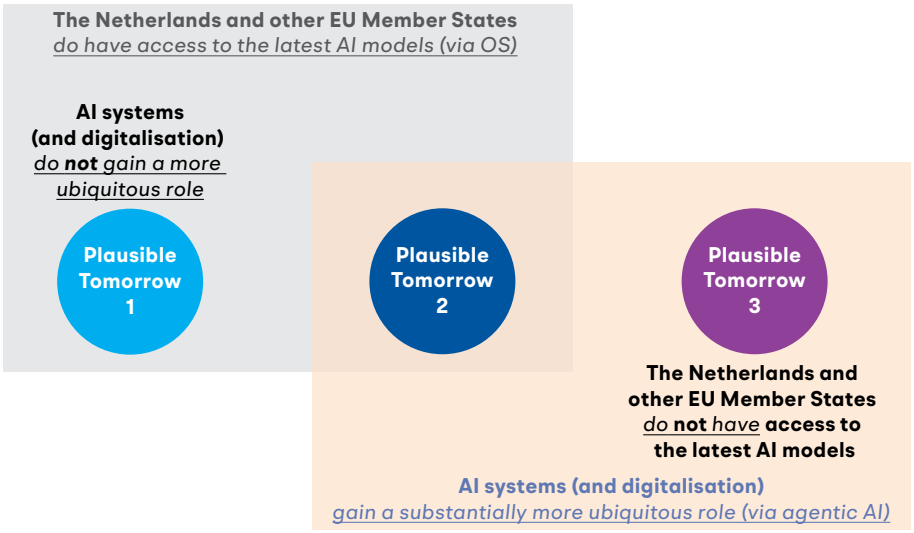

Scenario

Plausible Tomorrow 1: AI Lights Up the Port of Rotterdam

How the proliferation of (open-source) AI offers a power boost to non-state actors, opening a Pandora's box

CONTEXT

AI models have continued to develop in recent years, albeit at a slower pace than some had hoped. After Deep Seek released its open-source frontier model back in 2025, open-source became the dominant approach. **AI thereby became more widely accessible than ever before**: AI proliferation and widespread adoption are now a given. Not only states, but also non-state actors possess more substantial capabilities than ever before and are developing a wide range of AI applications for political, social, economic and military gain. This creates opportunities but also huge challenges.

Description

Given the rapid advances both in capabilities and availability of open-source AI systems, it is becoming easier for malicious actors to abuse AI. Self-navigating drone swarms, for instance, are commonplace weapons used by not only state actors such as Russia and Ukraine, but also by terrorists, activists and journalists. The complexity of police and intelligence work is sky-high: following the proliferation of AI, the internet has been poisoned with disinformation, making truth-finding nearly impossible, and reducing trust and increasing polarisation throughout society.

The upcoming elections for the European Parliament in 2029 are being manipulated on a massive scale, by various actors spreading 'deepfake' videos of politicians delivering fake speeches or behaving badly. Meta can still tell some deepfake images from real ones, and as the owner of Facebook, they are better positioned to analyse the origins of posts on their platform. However, Meta has been arbitrary in its willingness to lend a helping hand, loosely trending about the fluctuations in President Trump's mood, who is serving a third term following the Second – 'AI-induced' – Capitol Siege in 2028.

Against the backdrop of all this uncertainty and the enhanced capacity to cause chaos from the widespread availability of open-source AI, a wave of drones, flying in perfect formation under the command of an AI navigation algorithm that constantly re-routes them to evade jamming, strikes its target. This time, the drones hit five large container ships at the Port of Rotterdam. As the drones' Cobra 5[17] payloads detonate on impact, windows shatter and dockworkers shudder. Minutes later, under the sinister sound of sirens, the port lights up. The ships glow red beneath a thick layer of black smoke: hundreds of Chinese humanoid robots incinerated at the dock. 'Another attack', sighs a tired police researcher, 'but who's behind this one?'

In Denmark, meanwhile, Danish conservative Rasmussen is pleased and somewhat relieved. 'This is sure to get me the missing votes!', he hopes. Later that day, the Danish parliament will vote on extensive legislation aiming to increase surveillance. Massive monitoring operations using advanced AI-enabled cameras will keep everyone safe and sound, promises Rasmus. In his pocket, he carries a letter signed by a group of 1,984 industry advocates.

17 Cobra 5 explosives are illegal fireworks with the power of a hand grenade that can be bundled. They are used in the Dutch criminal circuit and often attached to houses, sometimes with deadly results.

National Security Implications for the Netherlands in the EU Context

In this Plausible Tomorrow, AI is essentially open and widely accessible through open-source models. The Netherlands and the EU have access to the latest AI technologies, just like any other country as well as companies, NGOs, journalists, activists, criminals, extremists and terrorists. Overall, non-state actors are rapidly gaining information and military capabilities, narrowing the gap with nation-states. This wide availability makes it very easy to spread disinformation or coordinate online troll armies, and to enact physical attacks such as the one envisioned above. The Rotterdam drone strike is just the latest case in a new era of 'semi-autonomous terrorism'.

While AI contributes to detecting cyber threats and vulnerabilities, it equally empowers more users to carry out cyberattacks. AI-generated images and videos are nearly indistinguishable from real content, further weakening traditional media, which were already struggling under the pressure of AI-driven competition. Deepfakes and disinformation are no longer sporadic but systemic – 'truth decay' is now a structural condition. This is more than just a media problem: it corrodes the foundations of Dutch democratic governance, and social cohesion is breaking down fast and evermore difficult to avert. Citizens feel disconnected from public discourse, overwhelmed by distrust and a growing sense of fear towards AI systems, speculating about espionage via consumer apps like TikTok or Instagram and feeling a lack of confidence in institutional communication. Even essential sectors like healthcare become compromised. Actors may unknowingly adopt Trojan Horse AI tools, exposing systems to surveillance and cyberattacks. AI-generated content fuels conspiracy-driven civic unrest, for instance around the legitimacy of democratic elections and immigration.

Economically, societies become more vulnerable. The scenario imagines hybrid threats that disrupt physical and digital infrastructure, inevitably impacting the economy. There are also opportunities: in this scenario, the Netherlands makes use of open-source AI to expand its lead in sectors like AgriTech and HealthTech, and to support developments in new high-tech niches where the country has a strong position, like quantum technologies. Other countries apply AI to their existing strengths and new fields too: notably, AI allows the Global South to bridge the knowledge gap.

In the military domain, states struggle to regulate the proliferation of autonomous weapons. There are no clear standards or control mechanisms. The threshold for pre-emptive and preventive strikes lowers, while lethal autonomous weapon systems spread beyond traditional armies. The primary threat is no longer dominance by superpowers, but rather the uncontrollable scale and diffusion of destructive technology. Beyond direct attacks, adversaries may also target the systems that monitor and control Dutch airspace. AI-enabled drone swarms could deliberately jam, spoof or saturate air-traffic control radars and communication channels, creating unsafe skies and jeopardising commercial aviation. This risk links closely to surveillance, as when monitoring systems themselves are manipulated or overloaded, the ability to distinguish genuine threats from noise is eroded. In effect, AI turns the very infrastructure designed to ensure safety into a vulnerability, compounding both territorial and societal security risks.

Besides immediate cyber-, information and military threats, widespread open-source proliferation also poses ecological risks: the uncontrolled spread of energy- and water-intensive AI models accelerates pressure on Dutch and European sustainability targets, with data centres and chip production directly impacting the living environment. Moreover, the absence of effective international rules to govern such diffusion undermines the Netherlands' capacity to uphold the international legal order, as norms on autonomous weapons, cyber operations and digital sovereignty remain fragmented or ignored.

Many of these challenges already exist today – but the capabilities will be far more accessible to a wider range of actors. For the Netherlands, this Plausible Tomorrow presents a strategic paradox: while AI boosts innovation in key sectors, it also erodes national resilience through asymmetric vulnerabilities.

Scenario

Plausible Tomorrow 2: A Fragile Power (Im)Balance

How a fast-changing AI landscape drives geopolitical volatility rivalry and uncertainty to a max

CONTEXT

Developments in and around AI have progressed in highly disruptive ways in the years up to 2030 – both in closed and open-source models. Major and rapid advances in agentic AI have enabled AI's integration across a wide range of applications and hardware, such as smart devices and robots. As a result, **AI is omnipresent** and plays an all-encompassing role in political, societal, economic and military spheres. The US and China dominate in closed military AI systems, while some countries, including the Netherlands and Gulf states, possess second-tier but competitive closed models.

Description

The team behind rAIdar[18] is shocked to see that their scale-up has made the national news overnight. Alphabet had launched a hostile takeover bid for its patents and intellectual property, prompting the Dutch government to panic and block the deal. Initially focused on using AI to improve radar systems for applications in commercial cars, the military dual-use opportunities led the Dutch Ministry of Defence to place several large orders with rAIdar last year. The move was welcomed by many, as it meant avoiding yet another source of dependency on American Big Tech and alternatives from the United Arab Emirates. No European government wants to repeat the mistake of the UK way back in 2015, when it let Google acquire DeepMind – the most promising AI lab in the world at the time.

rAIdar is not the only Dutch company that manages to compete on the global stage. Similarly, it is not unique in finding itself entrapped in a tense dynamic involving ambition, research, the global tech race and power politics. The Netherlands and the EU have long been indecisive and insecure in their industrial

18 As of the time of writing, rAIdar is a non-existent, fictitious company invented for the purpose of describing this Plausible Tomorrow.

and innovation policies. They failed to rescue the tie-up between ASML and Mistral, which seemed promising just a few years earlier. With the return on investment from Mistral's LLMs in the red and ASML in the shadow of China's lithography boom, the collapse of two European champions is imminent.

Meanwhile, the US is overconfident and unpredictable. China still overly depends on advanced chips from the West, which are subject to strict export bans. Several other countries, such as Saudi Arabia, have moderately successful AI in different fields, but they dominate none.

Then, unexpectedly, the balance shifts. The United Arab Emirates unveils a compact AI-enhanced directed-energy weapon that is capable of disabling incoming missiles at a fraction of current costs.[19] Within weeks, China adapts the design and transfers it covertly to Russia, shifting the still ongoing war in Ukraine: Ukrainian offensive drones and missiles suddenly become obsolete. Western air defences suddenly look outdated. Yet the fragility cuts both ways: a rapid European breakthrough – such as secure photonic communications integrated with rAIdar's stack – could flip the board just as quickly.

While some AI applications, such as rAIdar's models, still partly rely on relatively traditional machine-learning applications, there has also been a boom in the use of agentic AI applications. The Netherlands has secured a leadership position in HealthTech. At the first-aid department of Dutch hospitals, patient triage has been fully automated. The decision-making process for where to send respective patients has sped up, halving the average waiting time for initial diagnosis while uplifting accuracy rates to over 98 per cent. In addition, AI agents are quicker and more accurate at drafting reports and patient dossiers, allowing doctors to focus on healing patients instead of doing paperwork.

Not all is well with agentic AI, however. In the US, the administration has introduced automated AI judges in federal criminal courts. The government deemed this necessary to ensure a fair and efficient judicial branch, while moreover saving money on an expensive institution. Conversely, human rights activists point at leaked government documents that reveal that the administration has knowingly and purposefully used an underlying LLM that increases racial inequality and bias in court rulings. The 'black box' nature of AI systems allowed the government to do so under the guise of introducing 'an innovative and neutral computer program'.

19 Directed-energy weapons are systems that use focused energy – such as high-power lasers, microwaves or particle beams – to disable, damage or destroy targets without traditional projectiles.

National Security Implications for the Netherlands in the EU Context

In this Plausible Tomorrow, agentic AI has rapidly evolved, driving not only economic transformation but also geopolitical reconfiguration. Global complexity has surged beyond the capacity of most states to navigate technological disruption independently. Competition is fierce in all layers of the technology stack, now mostly on algorithms and applications. Rivalry continues to exist in the hardware layers because of decreasing access prices and innovative breakthroughs that help to elevate those with more limited resources.

Companies struggle to stay afloat in this highly competitive and rapidly evolving environment. A fear of missing out on tech leadership and economic competitiveness reinforces the strategic alignment between states and their tech champions, as both sides increasingly feel the need to protect and assist each other amid intensifying geopolitical rivalry.

Governance models fail to keep up with rapidly developing and commercialised AI technologies. Global coordination on AI ethics and safety standards collapses under the weight of divergent state interests and corporate influence. Regulation remains largely symbolic and fragmented at the national and regional levels, as all power blocs aim to gain an edge over others.

Meanwhile, the economic landscape is marked by massive disruption. AI-induced automation drives deep restructuring of the labour market. The widespread availability of open-source models and agentic AI makes many knowledge-based professionals superfluous, including junior lawyers, graphic designers and programmers. Deemed indispensable just a few years ago, many of them are now replaceable by armies of AI agents controlled by a small number of specialised AI prompters. This puts huge pressure on the social security and cohesion of European welfare states, and the debate around policies like the Universal Basic Income now takes centre stage in the Netherlands. As some studies predict that up to 50 per cent of jobs may disappear across white-collar sectors before 2035, questions about social protections and income models are central. With this wave of unemployment coming, governments face starker choices than ever: as most profits leave the EU, they either have to increase income tax, and so risk crippling consumption; or reduce expenditure on social welfare or healthcare, thus increasing social unrest.

Truth decay further accelerates in a world dominated by agentic AI. Autonomous agents can generate, personalise and distribute tailored propaganda at a scale beyond human moderation capacity.

The penetration of agentic AI into critical services also generates risks for physical safety in a whole number of domains: failures in automated triage or biased decision systems in healthcare, for instance, could endanger patients rather than protect them. At the same time, the climate impact of mass-scale AI and its necessary energy-intensive data centres puts pressure on the European Green agenda and electricity grids. EU Member States and regions race to build their own Silicon Valleys – often without sustainable energy strategies.

Militarily, the blurred boundary between private innovation and strategic weaponisation becomes a source of risk. Highly innovative tech companies benefit ever more from first-mover advantages, offering dual-use systems that create new threats – and then proprietary solutions. The number of companies following the surveillance-based business model of Palantir has surged and there is little democratic control over them. The proliferation of cheap, autonomous and destructive technologies outpaces legislation and control regimes.

Disruptive AI is no longer an anomaly – it is the default, and geopolitical blocks are fighting to outflank each other.

Scenario

Plausible Tomorrow 3: White Flag Festival: Surrender to the AI Gods

How the Netherlands and Europe are forced to obey either the Americans or the Chinese

CONTEXT

Developments in and around AI have been highly disruptive in the preceding years. AI models and applications are widely used by governments, companies and individuals, and in all economic sectors and elements of social life. AI now plays an indispensable role across political, societal, economic and military domains. The US and China have developed into unrivalled AI superpowers and compete for the supremacy of their systems in third countries and regions. Only the US and China possess highly advanced AI systems, leaving other countries – including the Netherlands – entirely dependent on one or both of these AI superpowers...

Description

Late 2025 saw Donald Trump pose an ultimatum to the EU and its Member States, which Trump viewed as a liberal project, out to destroy the US. The EU had a choice, he said, between Chinese tech or US tech. If the EU continued to allow 'undemocratic Chinese tech', thereby 'supporting the communists', Trump would cut off the EU's access to AI products from the US. Moreover, the EU would have to let go of most of its ethical AI regulations, as well as significantly reduce Chinese imports. With a digital knife at its throat and its back against a wall of technological overreliance, the EU felt it had no choice but to comply...

As American AI systems crept into European society, taking over jobs, Big Tech became Bigger Tech. It seemed so convenient at first, at least for the lucky ones: they could work less and spend more time with their kids or enjoying iced matchas in the sun. As the EU and EU Member States gave up on their short-lived ambition for (digital) strategic autonomy, the EU's digital deficit, like that of Canada, Japan and others, became an economy-wide trade deficit, while European capital in Silicon Valley skyrocketed. Mass layoffs of white-collar workers resulted in the complete rupture of social cohesion. Some politicians proposed special taxes for AI and social media products to offset the massive outflow of capital, but Trump blocked any

attempts to that end, threatening to retract the US from NATO. NATO chief Mark Rutte assured President Trump that all would be okay, stating that NATO countries would soon spend 5 per cent of GDP on defence and another 3.5 per cent on dual-use AI.

Now, in 2029, it has become clear – even to critics of the EU's large-scale adoption of American AI – that countries aligning with China were in an even more precarious position, at least for the moment. On 4 July 2029, ChatGPT 15o5 was released, focused on upgrading Tesla's humanoid robots to automate childcare in Western countries. As the demand for electricity in Europe surged with this upgrade, power gids started to fail, resulting in a national power outage in the Netherlands and several other European countries.

That same day, however, Russia's growing dependency on Chinese technology began to backfire. Beijing announced that it was unhappy with Moscow's floundering economy, and quietly imposed new conditions on continued access to its AI and robotics systems. Huawei and other Chinese tech giants rolled out software 'security updates' that effectively gave Beijing full operational oversight of Russia's automated industries and digital infrastructure.

When several regions resisted, important parts of their power grids and logistics networks were suddenly paralysed. Critical robots in logistics hubs and energy plants stopped responding to Russian commands, waiting instead for remote instructions from Beijing. With production halted and public order fraying, local authorities were forced to comply with Chinese directives, effectively turning Russia's economy into a controlled dependency. What had begun as a partnership of convenience had evolved into a digital vassalage.

It was day three of the power outage as Sytse, a millennial and proud early-adopter of new technologies and gadgets, let out a sigh of relief. With the shutdown of AI came a temporary relief from the 'AI colonialism' to which the Netherlands had been subjugated. True, it was painfully difficult for Sytse to relearn to plan routes, form opinions, make choices – in short, to live in an offline world. Still, as barter and trust replaced payments through plastic and bytes, Sytse felt in control and more human than he had in a long time.

The power outage brought some time to reflect truly, free from intrusive AI advice. Would a choice for European AI have been better? Could the EU have opted to advance without access to cutting-edge AI models? What would have been the hard and soft power implications of such choices? Given the lively barter economy that Sytse was now witnessing, perhaps the data centres could remain off a little longer, once the lights turned back on?

National Security Implications for the Netherlands in the EU Context

In this Plausible Tomorrow, the global AI race has hardened into a binary contest between the United States and China. Open-source models are widely available but significantly inferior in capability to the closed, proprietary models from the US and China. The rest of the world becomes secondary terrain, where sovereignty is traded for access to AI systems. The EU and its Member States no longer have control over their own digital society and economy – effectively turning them into digital colonies. The dream of strategic autonomy, societal resilience and economic security has collapsed.

The two superpowers now fully dominate the AI supply chain, from infrastructure (including data centres, cloud services and AI chips) to algorithms and applications. The preferences of American and Chinese tech giants increasingly shape domestic political decisions as well as diplomatic policies and actions. The world is geopolitically split in two. Countries either work for and obey America or China. In the absence of European-built AI propositions, the continent is vulnerable to exploitation and engineered instability: after all, it depends on systems designed elsewhere to control critical infrastructure and public functions.

Without urgent investment in AI infrastructure, education and manufacturing, the EU essentially becomes a vassal in a techno-feudal system defined by others.[20] The military implications are stark: EU and Member States' capabilities depend on second-tier technologies, leaving the continent in a dangerous position. Autonomous weapons, drone swarms and agentic AI-powered surveillance define modern warfare – but Europe lacks control of foundational systems. As US and Chinese companies race towards militarised AI, the EU can only respond with moral appeals in public, while begging for access to imports behind closed doors.

European states find themselves politically powerless. Media moderation, labelling and narrative construction are outsourced to AI systems that are tuned in either Beijing or Washington. As truth becomes centralised in the hands

20 *The Guardian*, '"Capitalism is Dead. Now We Have Something Much Worse": Yanis Varoufakis on Extremism, Starmer, and the Tyranny of Big Tech', 24 September 2023.

of a few private or authoritarian truth arbiters, the EU's capacity to maintain legitimacy in the digital public sphere is severely weakened.

In this dependency scenario, truth decay takes a new dimension: people are increasingly shaped by dependencies on foreign systems – with no meaningful regulatory oversight to control this. The Netherlands risks ceding not only digital sovereignty but also narrative sovereignty: citizens internalise framings of security, identity and norms that reflect outside interests rather than domestic democratic debate. Human oversight becomes decorative – a token presence amid sweeping automation.

Economically, Europe slips further behind. The continent faces capital flight, loss of tech-industry relevance and shrinking innovation ecosystems. Skilled labour moves abroad. Even ASML, once Europe's ace card, is buckling under foreign pressure. In the age of AI gods, there is little room for minions: the Netherlands is about to watch its most strategic company depart, leaving the country – and the continent – truly empty-handed. Lacking access to the latest AI innovations, European companies and countries face an increasingly uphill struggle. States may be forced to barter natural resources or regulatory compliance just to access critical systems.

Dependence on foreign AI systems also erodes protections for Dutch citizens' physical safety: with safety-critical updates controlled from abroad, the ability to safeguard life and wellbeing within the Netherlands diminishes. Finally, the effective outsourcing of governance to US and Chinese technology giants undermines the Netherlands' commitment to the rule of law, as compliance with treaties and human rights standards becomes subject to external political and corporate interests.

This Plausible Tomorrow underscores that tipping points matter: by the time Europe recognises its loss of relevance, the moment for effective action has already passed, leaving the bloc paralysed and dependent.

From Foresight to Execution: Policy Priorities for the Netherlands and the EU

Drawing on the plausible tomorrows above – and anchored in our baseline of ten key takeaways that underscore the strategic race between great powers, the unpredictability of generative AI timelines, the risks of proliferation and misuse, the vulnerabilities of digital dependencies and the ecological costs of scale – this section distils concrete lessons and identifies policy priorities. Table 1 summarises the estimated impact of each Plausible Tomorrow on the six Dutch national security interests detailed in the introduction.

Table 1 Impact of each Plausible Tomorrow per Dutch national security interest

National Security Area	Impact of Plausible Tomorrow 1	Impact of Plausible Tomorrow 2	Impact of Plausible Tomorrow 3
Territorial security	Medium	Medium	High
Physical safety	Medium	Medium	High
Economic security	Medium	Medium	High
Ecological security	Medium	High	High
Social and political stability	Medium	High	High
International legal order and stability	Low	High	High

Across all three Plausible Tomorrows, malevolent actors use AI to destabilise Europe's information environment, turning sporadic disinformation into a constant stream of alternative truths that polarise society and weaken democratic legitimacy. These pressures are compounded by vulnerabilities in critical infrastructure, economic resilience and ecological sustainability. While Plausible Tomorrow 1 brings medium risks from uncontrolled proliferation of open-source AI models across state and non-state actors, Plausible Tomorrow 2 amplifies disruption through autonomous, agentic systems (so-called agentic AI). Plausible Tomorrow 3 delivers severe impacts across all six security interests as Europe is essentially a digital colony of foreign powers.

One key takeaway is that the Netherlands and the EU must act swiftly and firmly in a variety of domains to continue protecting their societies from AI-induced threats and should seize opportunities to build leadership. A key element herein is to pair safeguards for ethical AI with an assertive industrial strategy that nurtures Dutch and European capabilities. To move into this direction the Netherlands and the EU could consider the following recommendations:

First, **restrict, regulate and license dual-use autonomous capabilities on EU soil before they cascade to malign non-state actors**. The Rotterdam drone strike vignette underlines how consumer tech can become weapons systems. The mysterious drone flights at several European airports in September 2025, as well as the drone attack by a terrorist group towards a police helicopter performing an anti-drugs operation in Colombia in August 2025, show that this scenario is far from science fiction.[21] It is crucial for the Netherlands and the EU to critically assess high-risk consumer autonomy systems (for example, advanced drones and humanoid robots), ensuring that licensing, registration, geofencing[22] and manufacturers' duty-of-care are non-negotiable and aligned with the EU AI Act's application risk tiers.[23] This also includes strengthening AI literacy obligations for manufacturers and preparing the Netherlands' competent authorities to enforce prohibitions and high-risk requirements as they come into force.

Second, **mandate safe update practices for autonomous machines**. The Plausible Tomorrows show how robots or self-driving systems could be compromised via malicious vendor updates. Independent audit and certification processes for software/firmware updates in any machine capable of autonomous movement are essential, so compromise cannot propagate silently at scale. This is no different than the road safety and food security checks that currently protect Dutch and European citizens from daily life risks.

Third, **safeguard the information environment**. Generative AI – not only text, but increasingly images and video as well – will continue to supercharge disinformation and truth decay, especially during electoral cycles. EU Member States are the ultimate bulwark to protect liberal democracies, and must behave accordingly. Scaling critical-thinking and digital-media literacy across

21 *BBC*, 'At Least 18 Killed and Dozens Injured in Separate Colombia Attacks', 22 August 2025.

22 Drone geofencing uses GPS signals and flight-control software to create virtual boundaries in the sky, creating, for example, no-fly zones.

23 European Commission, AI Act, Regulation (EU) 2024/1689.

all education levels, the civil service and society at large must be a top priority, implementing an important pillar of the vision on generative AI proposed by the Dutch government in 2024.[24] At the European level, it is important for the EU to hold the line on implementing the Digital Services Act (DSA), resisting pressures from (mostly American) tech companies and the US government.[25] Instead, the DSA should be used as a blueprint and best practice even for companies outside of the Act's reach. The same applies for the AI Act. The 2025 European Commission's guidelines clarify the definition of 'systemic risk' for general-purpose AI models and the mechanisms whereby models can be classified as such, which is a step in the right direction.[26] Yet, as with the entire digital-related body of law that was produced during the 2019–2024 Commission tenure, actual implementation will define the measure of success. Calls during the second semester of 2025 by Mario Draghi, industry groups and even Swedish Prime Minister Ulf Kristersson to pause implementation of the AI Act, for example, should be met with strong opposition.[27]

Moreover, in the context of a perceived global AI arms race between the United States and China, Europe cannot afford to focus on regulation alone, alongside a reactive, hesitant posture overall. Safeguarding the information environment will only be credible if regulatory enforcement is matched by **continuous monitoring, rapid-response mechanisms for disinformation surges and cross-border coordination with allies**. This means coupling long-term literacy and regulatory initiatives with capacities such as national rapid-alert systems and public–private fact-checking networks. Such investments in societal resilience will enable the Netherlands and the EU to adapt at the same pace that malign actors innovate.

24 Government of the Netherlands, 'Dutch Government Presents Vision on Generative AI', 18 January 2024.

25 *Brussels Watch*, 'Trump Administration Intensifies Lobbying Against EU Digital Services Act', 7 August 2025.

26 The classification mechanism works both ways: owners of models that pass the threshold but hold no systemic risk can apply for exceptions (false positives); and the Commission can qualify models that do not pass the threshold but do present systemic risk (false negatives). See: European Commission, 'Guidelines for Providers of General-Purpose AI Models', 31 July 2025.

27 See: *Euractiv*, 'Draghi Calls for Deep Cuts to Privacy Rules and Pause on "High Risk" AI Act', 16 September 2025; *Le Monde*, 'IA : 45 entreprises européennes demandent une « pause » dans l'application de l'AI Act' (in French), 4 July 2025; and *EU Insider*, 'Swedish PM Calls for Pause on EU AI Rules, Citing Lack of Common Standards', 23 June 2025.

Fourth, **prevent algorithmic harms before they institutionalise**. (Agentic) AI systems are prone to incorporate bias. In critical sectors where citizens' sensitive data are managed, such as healthcare, justice or welfare, this risk can compound existing institutional shortcomings. The Dutch childcare benefits scandal of 2022 showed how damaging the combination of biased algorithms and lack of humans in the loop can be: the problem was not only the flawed risk model, but also a lack of effective oversight.[28] It is important that Dutch public bodies and vendors prioritise this concern, ensuring that the procurement of software or autonomous systems is contingent on demonstrated compliance with these principles. Where fairness cannot be assured, such as in the case of machine-learning algorithms, human oversight and non-automation serve as the appropriate default. This aligns with the Netherlands' human-centred, risk-based stance on AI and with EU laws' safeguards around high-risk usage of AI.

Fifth, a **'rights-based' approach** ought to be brought to the forefront of the public discussion.[29] Better understanding is needed of the AI Act's risk-based approach as a step towards a less dystopic situation than most doom scenarios propose. After all, the AI Act covers only a subset of an increasingly automated economy and society, contrary to what more than a few companies, government officials and critics seem to believe. In the Netherlands, as in many other countries, for instance, supermarkets have increasingly stopped employing humans in the check-out and payment area; and more and more banks are reducing the number of physical locations and counters that they offer. This touches upon the broader themes of digital divides, literacy and social exclusion. A rights-based approach would mean, for instance, that people who want to be served by fellow citizens should have the right to do so. Whether concerning the right to receive customer service from a human, the right to a human judge in court or the right to a human manager in a large warehouse, there are many cases where AI adoption should be strictly avoided.

Sixth, and crucially, **ensure long-term, sustained and deep investments in the competitiveness and indispensability of strategic industries** – the best possible defence across all Plausible Tomorrows. This means **building and adopting own capacity** by, among others:

28 *Politico*, 'Dutch Scandal Serves as a Warning for Europe over Risks of Using Algorithms', 29 March 2022.

29 Inspired by Dan McQuillan, 'EU AI Act Briefing', 9 December 2023.

- **Making the most of existing and planned domestic compute and facilities.** The upcoming AI factory in Groningen should be embedded in a proper Dutch national and European plan that connects present and planned facilities, ensuring that researchers, start-ups, small and medium-sized enterprises (SMEs) and public agencies can access and use capacity under their own terms. This approach leverages existing and incoming investments to reduce external dependency (as flagged in Plausible Tomorrow 3) and aligns with the EU's broader AI Continent Action Plan to turn European strengths into practical advantages.[30]
- **Support Dutch deep-tech supply chains.** Dutch national funding and EU co-financing are best channelled towards strategic domains where the Netherlands is already strong, such as semiconductors, robotics, and cloud and compute niches anchored in the Brainport Eindhoven high-tech ecosystem. Programmes ought to prioritise high-end scientific research and its translation into products, services or applications that deliver broad societal benefits and/or strengthen Europe's control over critical technological chokepoints. Prioritise high-impact sectors such as AgriTech, HealthTech and EdTech, which can bring direct benefits to the population. These strengths should be nurtured at home and promoted abroad as a way of further building market share, competitiveness and the standard-setting power of Dutch companies in third countries.[31] Public procurement should be leveraged to pool demand for key technologies, such as AI chips and cloud services, across EU Member States and to favour European solutions, thus ensuring that public spending strengthens domestic and European capabilities rather than flowing to non-European providers.
- **Anchor these policies within the Netherlands' broader digital strategy and national security framework**. The Dutch Digitalisation Strategy, updated in September 2025,[32] alongside the wider national Tech and Digital agendas, sets out the key pillars to operationalise. Expanding the pool of digital experts and channelling them to where it matters the most are essential:

30 European Commission, 'Shaping Europe's Leadership in Artificial Intelligence with the AI Continent Action Plan', 9 April 2025.

31 For more on this, see Alexandre Gomes and Maaike Okano-Heijmans, 'Connecting the Dots: Linking Digital Global Gateway to Local Sector-specific Needs', April 2025; and 'Investing in Trusted AI Partnerships: Why Europe Needs to Ensure Alternatives to State-linked and Proprietary AI Systems' (forthcoming Clingendael Policy Brief, autumn 2025).

32 Government of the Netherlands, 'The Netherlands' Digitalisation Strategy (NDS)', 2 September 2025.

strengthening infrastructure, driving SME digital transformation, enhancing cybersecurity and focusing on Dutch strengths and niche sectors. Secure data-sharing remains a core enabler, but moving towards shared, well-governed data infrastructures – or data commons – can and should embed trust and sovereignty by design.

- **Channel EU capital effectively.** To ensure good use by European researchers and start-ups of facilities such as AI factories, tap into the EU's InvestEU programme, the InvestAI initiative and forthcoming instruments under the AI Continent Action Plan to crowd in private finance. At the national level, the Dutch government could incentivise pension funds to adopt mandates that allow them to (co-)invest in European digital strengths and autonomy. Aligning with Brussels' drive to mobilise major funding for, for example, European AI factories and champions, Dutch programmes should be structured to qualify quickly for these EU envelopes. Positioning the Netherlands to optimise their usage should go hand in hand with increasing Dutch tech sovereignty. This also means improving structural conditions – such as fiscal incentives for companies that reinvest their profits into R&D&I or talent-retention policies – to build a self-sustaining ecosystem in the long term. An ambitious Netherlands would work towards lifting R&D spending to at least 3 per cent of GDP. At 2.23 per cent of GDP as of 2023, Dutch R&D spending is above the EU average, but below the OECD average and key peers such as South Korea, Japan, Germany and the United States, all of which consistently spend above 3 per cent of GDP on R&D.[33]

Seventh, and finally, **governance must match speed with capabilities**. To develop and implement a clear vision for how AI should serve Dutch society, the Netherlands could appoint a **coordinating Minister of Digital Affairs under the Prime Minister's Office**, resembling models adopted in Japan and the United States. This office's mandate would not only be to maintain a **permanent foresight function** – anticipating technological trends and risks – but also to articulate a long-term, values-driven vision for AI that aligns national security, industrial and digital policy agendas. Staffed with a lean office consisting of one or two officials from all the other ministries, this coordinating minister could also publish annual technology and AI readiness updates, tracking compute capacity,

33 Rathenau Instituut, 'R&D Investments in International Perspective', 26 June 2024.

skills, procurement and safety incidents in the country. In addition, the office could monitor sectoral and industry developments to identify the Netherlands' AI niches.

More fundamentally, such a governance structure should enable the Netherlands to confront a **core strategic question: whether to shift from treating AI largely as a private good towards recognising it as a public one.** That decision will shape how public capital is deployed – whether to build and sustain AI as a public good, developing common, shared infrastructures and digital commons, or to shape markets and standards that steer private innovation.

An integrated, capability-building posture grounded in this vision and approach would strengthen the Netherlands' position as a producer of trusted and competitive AI – resilient in the chaotic world of Plausible Tomorrow 1, adaptive in the arms race of Plausible Tomorrow 2, and less dependent on the power politics of Plausible Tomorrow 3. The Dutch government would do well to **bolster the digital commons** — shared, open and well-governed digital infrastructures that serve the public interest —, **strengthen the technological stack and turn its value-driven approach into an industrial advantage**, so that whichever tomorrow arrives, Dutch society and Europe at large are better prepared.

Appendix. Scenarios for the Next Three to Five Years: Strategic Foresight Methodology

How are today's and tomorrow's developments in AI going to have an impact on the Netherlands? The honest answer is: nobody knows – including (future) AI itself. Why? Because the **future is fundamentally unknowable**. This is true in a general sense and even more true when it comes to developments that are taking place at a vast pace and that are having a disruptive impact in nearly all thinkable domains. AI is such a rapidly evolving, high-impact field, affecting the political, economic, military and societal realms. This makes **forecasting impossible as well as detrimental.** To allow for sound policymaking despite that, we may apply foresighting.

Strategic foresight offers us the possibility to evaluate not one but **several perspectives of possible or Plausible Tomorrows** that could reasonably materialise. In other words, it is not about what will happen (forecasting) but **what could happen**. This study offers three scenarios that can help stretch our thinking in terms of what could be on the AI horizon and to offer perspectives on what that could mean for national security. If we want a Dutch AI and national security strategy and policy that are future proof, we need a strategy and policy that are **futures proof.**

The three Plausible Tomorrows are derived from expert discussions in a scenario workshop hosted at the Clingendael Institute in May 2025, with participants from think tanks, the Dutch government and the AI sector. The aim was to imagine the national security implications of four different scenarios across two axes of fundamental uncertainties. The four resulting 'quadrants', as we called them, were supplemented with a third dimension: the degree of availability of open-source AI models. The three scenarios that differ most from today's situation are the three Plausible Tomorrows presented in this report.

Two Uncertainties, Four Quadrants

Participants in the workshop were presented with a Dutch version of the following visual representation of the resulting quadrants:

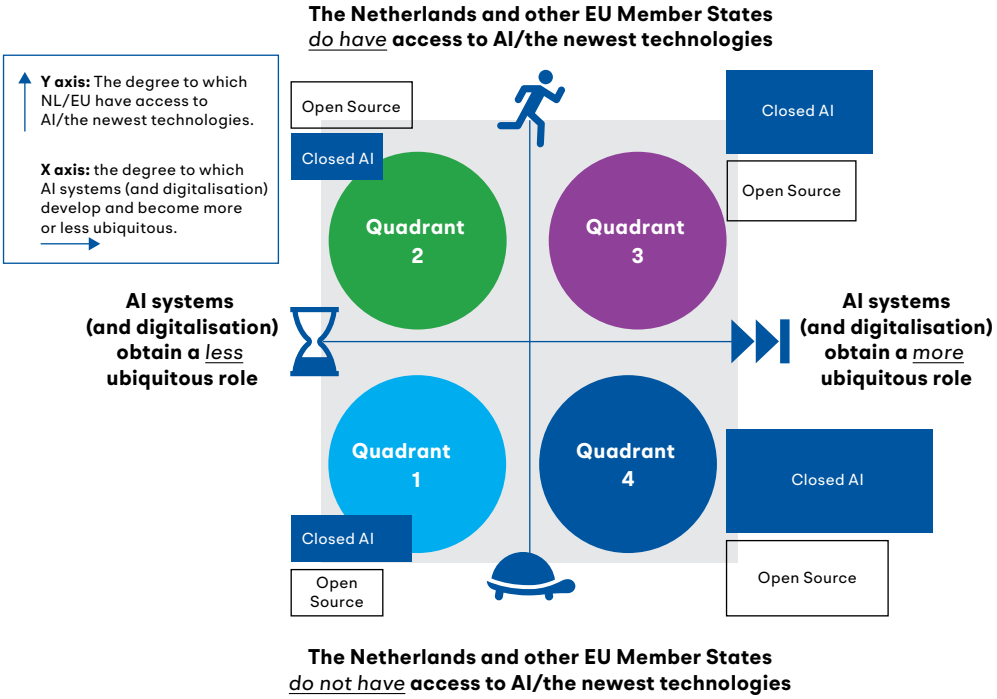

The four quadrants were summarised as follows:

1. **Quadrant 1: The Netherlands misses the 'AI boat' and lies by the wayside, while others prosper.** In this conceivable future, developments within and around AI have been less disruptive over the past few years. The Netherlands, like other EU Member States, has (very) limited access to AI and the very latest technologies arising from it. Open-source AI has developed to a limited extent. Closed-source AI systems, especially from US and Chinese companies, dominate the market. Of the four conceivable futures in this study, the one described in quadrant 1 most closely resembles today's status quo.
2. **Quadrant 2: Because of open source, everyone goes 'before the wind': to what extent is the Dutch ship really seaworthy for such a scenario?** In this conceivable future, open-source models have continued to develop, albeit at a limited pace by AI standards. Open source has thus become

dominant. This has made AI, as well as important AI applications, more widely accessible than ever before: not only state actors but also non-state actors have substantial capabilities and can thus develop a wide variety of AI applications, whether for political, social, economic or military purposes. In short, the Netherlands has the capabilities, as does (potentially) everyone else.

3. **Quadrant 3: 'The Netherlands is sailing, but in very turbulent waters'; (rapid/ disruptive development of AI and more access of AI).** In this conceivable future, developments within and around AI have developed very disruptively over the past few years. This is the case with respect to closed AI, as well as open-source models. In addition, tremendous progress has been made in the development of agentic AI, allowing this technology to be used in all kinds of applications and hardware (such as smart devices). As a result, AI has come to play a very comprehensive role on the political, social, economic and military stage. The US and China occupy a dominant position where closed military AI systems are concerned. Europe, including the Netherlands, the Gulf region and countries such as Brazil and Japan also have closed AI models that (can) be competitive. At the same time, companies and governments are struggling to keep their heads above water in this highly competitive and rapidly changing world. Widely available open-source models and agentic AI are adding to this and are having an impact in several sectors, such as education, healthcare, policing and for many knowledge workers.
4. **Quadrant 4: 'The Netherlands is at the mercy of the AI (weather) gods'.** In this conceivable future, developments within and around AI have developed very disruptively over the past few years. AI has come to play a very comprehensive role on the political, social, economic and military stage. Only the US and China have highly advanced AI. Other countries, including the Netherlands, are entirely dependent on one (or both) of these (AI) superpowers in this regard. Widely available open-source models are inferior to the closed models of the US and China in terms of capabilities. Agentic AI is vastly developed in this scenario, with far-reaching consequences, such as the widespread deployment of humanoid robots.