

AI-Driven Election Interference in 2025 Dutch Elections

The rapid evolution of artificial intelligence (AI) technologies has shown great promise for society at large, offering potential efficiency gains and scientific breakthroughs across nearly every facet of our lives. But the harms that come along with such an advancement in AI capabilities present an equally significant challenge to critical democratic infrastructure in countries like the Netherlands, which saw unprecedented attempts at election interference in its 2025 elections. A meteoric rise in AI-driven content online shows just how susceptible even the most highly trusted governments are to interference in their democratic processes because of AI deepfakes, algorithmic manipulation, and unchecked LLM interactions with would-be Dutch voters. In the case of the Netherlands, however, there are unique aspects of its coalition-style government and philosophy on free speech that could prove crucial to how the spread of AI-driven content is handled in an EU member country versus more polarized countries like the United States. To appreciate why this is the case, it's necessary to revisit the facts surrounding the AI-driven election interference in the Netherlands' 2025 elections.

The circulation of AI-generated content in the 2025 Dutch elections was novel due to the volume and scale at which it permeated online platforms and popular Dutch social media. A Netherlands-based consortium made up of academics, civil society, and digital forensics investigators, called the *Hybrid Election Integrity Observatory* (HEIO), released a report shortly after the 2025 elections that brings to light exactly how impactful AI-generated content can be in high-stakes democratic practices like an election. According to the HEIO report, while the elections remained free and fair, they “took place under unprecedented digital pressure that tested the resilience of democratic institutions and exposed vulnerabilities in the digital information ecosystem.”¹ This report and its findings are important for several reasons, but most

¹ HEIO Consortium. *HEIO - Dutch General Elections Oct '25 - Final Report*. 2026. <https://www.heio.nl/wp-content/uploads/2026/01/260115-HEIO-Final-Report.pdf>.

notably it highlights the potential threat to democracy that AI and AI-generated content poses, and it suggests that new methods of AI governance will be needed to protect democratic institutions like voting from further degradation as AI evolves.

To understand these threats and potential solutions, it's necessary to understand exactly what was occurring online and in social media leading up to the 2025 Dutch elections, as well as the types of AI-generated content that were being produced. More specifically, the AI-generated content that was used to influence the elections can be broken down into three primary components: synthetic media, social media and algorithmic manipulation, and large language models (LLMs) or chatbots.

Synthetic media, such as deepfakes, dominated much of the AI-generated content that was circulating online at the time of the elections. MIT defines deepfakes as “a specific kind of synthetic media where a person in an image or video is swapped with another person's likeness.”² According to the HEIO report, synthetic media like deepfakes became more extreme during the elections. The report details a couple of examples, including an AI-generated protest song, “Wij zeggen nee, nee, nee tegen een AZC,” which roughly translates to “We say no, no, no to an asylum seekers' center.” The song was created by an independent actor associated with the far right, and the song itself describes the Netherlands as being “full” and “enough is enough,” in reference to asylum seekers and immigrants. The song reached the top of the charts in the Netherlands before it was eventually removed by Spotify and YouTube.³ The same HEIO report found that this song went on to spawn thousands of TikTok videos over the course of several weeks, further spreading a message of hate that was intended to fear-monger

² Somers, Meredith. “Deepfakes, Explained.” *MIT Sloan School of Management: Ideas Made To Matter*, July 21, 2020. <https://mitsloan.mit.edu/ideas-made-to-matter/deepfakes-explained>.

³ NL Times. “Dutch Extreme Right Anti-Refugee Anthem Removed from Spotify, YouTube.” *NL Times*, November 12, 2025. <https://nltimes.nl/2025/11/12/dutch-extreme-right-anti-refugee-anthem-removed-spotify-youtube>.

potential voters ahead of the elections. This is only one example of the types of AI-generated content that were spreading rapidly across Dutch society (online) at the time.

Another well-known example of AI-generated content was a Facebook page that was among the most popular political pages in the Netherlands at the time. The Facebook page was titled “We will not press charges against Geert Wilders” and it circulated an unbelievable amount of AI-generated images and videos targeting immigrants and other minorities that the far-right party, Dutch Party for Freedom (PVV), has long attacked.⁴ The *EU Observer* reported that the page received nearly 75 million views between June and October of 2025, and the majority of the content on the page followed a common theme of pitting Dutch citizens against immigrants in a fight for resources and economic prosperity. The especially concerning part about content like this circulating on Facebook, or far-right songs reaching the top of the charts on Spotify, is just how easy it has become for bad actors to gain visibility. The World Economic Forum’s 2026 Global Risks Report also highlights this type of AI content generation as a top global risk. It states that in just the last five years alone, deepfake creation has become cheaper and more convincing, and as a result, it’s starting to have a greater influence on politics and election cycles.⁵ The same Global Risk report explicitly calls out the Netherlands, along with the U.S., Ireland, Japan, Argentina, and others as countries that have had to contend with this challenge in recent elections. This highlights just how quickly deepfake content is evolving and how far-reaching the implications will be on governments around the world, not just the Netherlands.

In addition to synthetic media like deepfakes, the HEIO report also highlights social media and algorithmic manipulation as yet another form of interaction with AI content that attempts to interfere with individuals’ political beliefs. The report details the coordinated use of

⁴ Carpenter-Zehe, Olivia. “AI-Driven Digital Manipulation ‘tested’ Dutch Election Integrity, Researchers Warn.” January 16, 2026. <https://euobserver.com/197303/ai-driven-digital-manipulation-tested-dutch-election-integrity-researchers-warn/>.

⁵ World Economic Forum. *The Global Risks Report 2026*, 21st ed. (Geneva: World Economic Forum, January 14, 2026), https://reports.weforum.org/docs/WEF_Global_Risks_Report_2026.pdf.

social media interactions such as “likes,” reposts, and attempts to make fringe content seem mainstream. A partner to the HEIO consortium, Trollrensics, identified a network of at least 23,000 accounts that liked the Facebook page for a left-leaning party leader, Frans Timmermans, while on X (formerly Twitter), thousands of accounts from Nigeria, Ghana, and Ivory Coast retweeted far-right political content.⁶ This sort of interaction with fringe political content on social media was not isolated. In fact, the Netherlands-based Rathenau Instituut, a Dutch watchdog agency, issued a report about this type of interaction (algorithmic manipulation and spread of misinformation by trolls during elections), warning about how widespread it was.⁷ The same institute went on to recommend better monitoring of online platforms by government institutions.

The final type of AI-generated content that was highlighted in the HEIO report was in reference to LLMs, or chatbots. According to Reuters, the Dutch Data Protection Authority (DDPA) warned voters ahead of election day that popular LLMs were pushing people to the political fringes. It reported that “While 15 different parties hold seats in the 150-member Dutch parliament at present, chatbots told voters in 56% of cases to choose between the far-right Freedom Party or the Labour-Green Left coalition.”⁸ The same reporting went on to state that even in instances where the DDPA tested feeding the LLM a specific campaign or party platform associated with less popular coalitions, the LLM continued to recommend only 1 or 2 of the most popular Dutch political parties to chatbot users.

In case the potential harms of AI-generated content that were uncovered during the 2025 Dutch elections are not obvious enough already, the concern goes beyond the content

⁶ HEIO Consortium. *HEIO - Dutch General Elections Oct '25 - Final Report*. 2026. <https://www.heio.nl/wp-content/uploads/2026/01/260115-HEIO-Final-Report.pdf>.

⁷ DutchNews.nl. “Beware the Trolls, Institute Tells Government as Elections Near.” March 11, 2026. <https://www.dutchnews.nl/2026/03/beware-the-trolls-institute-tells-government-as-elections-near/>.

⁸ Parodi, Alessandro, and Toby Sterling. “Dutch Watchdog Warns Voters against Using AI Chatbots Ahead of Election.” *Reuters*, October 21, 2025. <https://www.reuters.com/technology/dutch-watchdog-warns-voters-against-using-ai-chatbots-ahead-election-2025-10-21/>.

and underlying AI technology itself. From an ethics perspective, there are broadly two levels of concern when it comes to AI-generated content on democratic societies: first, the erosion of democratic institutions, and second, the loss of voter autonomy. In both cases, there is a common thread: the undermining of the European-style coalition government. When AI-generated content and other forms of mis- or dis-information are circulating heavily online, it often gives voters the perception that they have fewer choices available to them than there are. In the Netherlands, there are over a dozen political parties represented in parliament, many of them small-medium sized groups that will join coalitions with the most mainstream parties to form a government. The process of forming a government is complicated, and bargaining between political parties can sometimes take months.⁹ With the rise of AI-generated content and associated misinformation campaigns, the coalition style of government-building in the Netherlands, like other European governments, risks being stratified into fewer political buckets. When voters are pushed towards far-right or far-left political parties by LLMs or algorithmic manipulation on social media, there are less opportunities for compromise by way of coalition-forming among the smaller, less popular political parties. Interestingly, coalition-style democratic governments have been found to be less hospitable to the polarization of extreme identities and political ideologies.¹⁰ According to Protect Democracy, this is, in part, because they bring a broader spectrum of voices together over time and require compromise among those groups. This, in turn, means politicians find success in looking for common ground, rather than attempting to appeal to the far ends of the political spectrum, as is the case in two-party systems like the United States. In terms of LLMs or algorithmic manipulation, coalition-style governments like the Netherlands are perhaps being mutated, by extension of AI's pushing their

⁹ Netherlands Institute for Multiparty Democracy & Instituut voor Publiek en Politiek. "The Dutch Political System in a Nutshell." 2008. <https://nimd.org/wp-content/uploads/2025/02/Dutch-Political-System.pdf>.

¹⁰ Tudor, Grant. "Polarization, Extremism, and Proportional Representation." September 11, 2023. <https://protectdemocracy.org/work/polarization-extremism-and-proportional-representation/>.

constituents to just one or two parties, into a political environment with less consideration to the “in-between” parties.

This impact on the Dutch coalition-building form of governance is not unique to the Netherlands. In Germany, very similar dynamics were at play during its 2025 elections. A report by the Institute for Strategic Dialogue (ISD) in Germany found that similar attempts were made by both foreign (Russia) and domestic (far right) groups to exploit societal divisions in Germany around contentious topics like immigration, the war in Ukraine, and more.¹¹ Similarly, these actors leveraged AI-generated content to run sophisticated misinformation campaigns across social media and various online outlets. The same report details how German election oversight committees and democratic nonprofits have called for increased transparency and regulatory reform to strengthen democratic resilience, much like the Dutch and the HEIO.

The challenges that European governments are facing when it comes to AI-related interference in elections and democratic institutions—except for AI technology itself—are not necessarily new, but the scale at which it’s being deployed is. Part of the success of bad actors in driving this interference is due to what the Rathenau Instituut refers to as “negligent influence.” In its report, “Scrolling to the ballot box,” the Rathenau Instituut describes the impact of emotionally charged content online. Research suggests that content is far more likely to go viral if it evokes a specific emotion.¹² This, in turn, leads to more amplification of, and engagement with, content that elicits strong reactions around topics like immigration, war, the economy, and so on. The same report finds that, because of this phenomenon, recommendation algorithms distribute these divisive messages to more people more quickly.

¹¹ Institute for Strategic Dialogue, Alliance4Europe, Debunk.org, GMF Alliance for Securing Democracy, DEN Institute, & EU DisinfoLab. *Foreign Information Manipulation in the 2025 German Federal Election*. Alliance4Europe, 2025. <https://alliance4europe.eu/foreign-information-manipulation-2025-german-federal-election>.

¹² Ex, Luuk, Quirine Van Eeden, Wouter Nieuwenhuizen, and Mariette Van Huijstee. *Scrollend Naar de Stembus: De Rol van Aanbevelingsalgoritmen Bij Inmenging in Verkiezingen*. Rathenau Instituut, 2026. [https://www.rathenau.nl/sites/default/files/2026-03/Scrollend naar de stembus Rathenau Instituut.pdf](https://www.rathenau.nl/sites/default/files/2026-03/Scrollend%20naar%20de%20stembus%20Rathenau%20Instituut.pdf).

Further, platforms are now using AI to make assumptions about what voters want based on what content they engage with online—even if it was content that was essentially force-fed to them by the algorithm.

This presents a major problem for voters: as more AI-generated content circulates online and is fed to populations who aren't necessarily asking for it, voter autonomy is quickly degraded. According to the University of Amsterdam (UvA), this ethical dilemma is further exacerbated using AI technology like LLMs. UvA found that 1 in 10 Dutch voters are likely to ask AI chatbots for election advice,¹³ which is often not reflective of their actual preferences, as mentioned previously. The UvA report confirmed that even in cases where voters' positions might have perfectly matched that of a middle-ground or smaller party, most often the chatbot recommended one of two major parties—regardless of what the user might have explicitly asked or searched for.

The erosion of democratic institutions and loss of voter autonomy may be the biggest ethical dilemmas at the forefront of the challenges associated with AI-generated content, but they're not the only ones. As is often the case with complex, multifaceted attacks on democracy, even seemingly “good” solutions can create new harms. In this case, there are several instances where gaps within existing governance mechanisms have resulted in new problems. A clear example of this is related to the restrictions on political ads leading up to elections in the Netherlands (and the EU as a whole). According to the HEIO report that investigated the 2025 Dutch elections, bans on political ads that were adopted by Meta and Google (as a result of requirements by the EU for political ad transparency) often pushed voters to other, less moderated parts of the internet.¹⁴ As a result, there was even less oversight than if the ads had

¹³ de León Williams, Ernesto, and Fabio Votta. “1 in 10 Dutch Citizens Are Likely to Ask AI for Election Advice. This Is Why They Shouldn't.” *The University of Amsterdam*, October 29, 2025. <https://www.uva.nl/en/shared-content/faculteiten/en/faculteit-der-maatschappij-en-gedragswetenschappen/news/2025/10/dont-ask-ai-for-election-advice.html>.

¹⁴ HEIO Consortium. *HEIO - Dutch General Elections Oct '25 - Final Report*. 2026. <https://www.heio.nl/wp-content/uploads/2026/01/260115-HEIO-Final-Report.pdf>.

been hosted on a major platform to begin with. Overall, the collective harms of AI-generated content and its cascading effects on the spread of misinformation and inadequate content moderation created a collective digital pressure that pushed the Dutch election system to its limits. While the HEIO report holds that the elections remained free and fair, they were not free of bad actors and political harms that shaped the online discourse at the time.

Having said this, the Netherlands is a country that enjoys high trust in government by its citizens and is uniquely capable of a response that can do these harms justice. And because of real differences in their legal system, such as the lack of a direct first amendment equivalent like we have in the U.S., there may be more leeway for them to leverage in the fight against AI-driven interference in democratic processes. Unlike the U.S., the European Union goes further in its restrictions on free speech, prioritizing space for the handling of hate speech, public safety, and national security.¹⁵ The EU's legal philosophy and priorities are different, and in the case of fighting harmful AI-driven content, likely more beneficial. Going back to the examples of the far right hate song that popped up leading to the 2025 Dutch election, this is something that the Dutch legal system was able to get removed without much controversy. In the U.S., a takedown request for a far-right song online would likely face much more criticism by free speech advocates. In this way, the Netherlands (and other EU countries) are better situated both philosophically and legally to handle the removal of AI-driven content that causes harm.

As AI-driven election interference proliferates across the globe, it's critical that governments step in to curb the spread and impact of this technology on vulnerable democratic institutions. And while there's no "silver bullet" to combat this challenge in one way, there is consensus among organizations like the HEIO in the Netherlands that a multi-pronged strategy is necessary to combat it. In fact, the HEIO itself is a great example of a collaborative "task

¹⁵ Car, Polona, and Beatrix Immenkamp. *Hate Speech: Comparing the US and EU Approaches*. PE 772.890. European Parliamentary Research Service (EPRS), European Parliament, 2025. [https://www.europarl.europa.eu/RegData/etudes/BRIE/2025/772890/EPRS_BRI\(2025\)772890_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2025/772890/EPRS_BRI(2025)772890_EN.pdf).

force” of sorts that is supported by the Dutch government, making their recommendations more possible to put into action.

In particular, the HEIO report on the 2025 Dutch elections makes a broad call to restore voter autonomy and strengthen democratic resilience via four main channels: creation of election integrity task forces, requirements for algorithmic accountability, enhanced misinformation tracking, and human-in-the-loop safeguards. In the Netherlands, the HEIO itself has already served as an invaluable task force by highlighting the harms of AI-generated content during election cycles and proposing concrete solutions. Specifically, their report calls for a permanent, state-funded institution that can best tackle the oversight of elections and potential interference going forward. The HEIO suggests that real time monitoring to catch “algorithmic steering” and the implementation of a 24 hour take down request for deepfakes will be two potential solutions to mitigate harm leading up to election day. The HEIO’s second suggestion, relating to algorithmic accountability, suggests that platforms should offer non-algorithmic feed options, as well as 3rd party auditing by researchers and civil society. In addition to these solutions, the HEIO also calls for better ways to track the spread of misinformation on “dark social,” or under moderated platforms. This, in conjunction with more humans-in-the-loop and AI labeling, the HEIO proposes, will help to cut down on the overall quantity of AI-generated content on the web and as a result, lessen the impact.

But the HEIO’s recommendations are just those—they have not been proven out in the Netherlands specifically. For this reason, it is worth considering the success of other democratic countries in fighting the same challenges, and perhaps no nation has done more to tackle this than Taiwan. As a country that has been vigilant against authoritarianism for decades, they’re particularly sensitive to any form of interference in their online political discourse. According to a Global Taiwan Institute (GTI) overview, they passed broad legislation in 2023 around the handling of deepfakes and established a task force to carry out their “4C strategy,” a framework

to cut, counter, cure, and collaborate against misinformation and AI-generated content.¹⁶ The first part of their strategy, cutting production of AI-generated content, calls for deterrence mechanisms that make the spreading of such content less enticing for potential bad actors, and increase the costs of doing so. The same GTI overview states that the Taiwanese government has increased criminal liabilities for these cases to the maximum sentencing time and established well-funded task forces to more quickly trace the origin of AI-generated content and misinformation online. The second strategic approach, clarifying falsehoods, involves fact checking with speed. The GTI overview confirms that Taiwan implemented a “222” approach, which “requires every government agency to debunk misinformation within two hours using two images and 200 words, crafted to be especially sharable and understandable on social media.” The third strategy of their 4C framework, curbing dissemination, relies on broadcasters and independent media to immediately remove any deepfake content upon verification of its inauthenticity. And the last strategy outlined in the GTI overview is referred to as “cultivating digital literacy,” which leans on campaigns run by civil society and nonprofits to inform the populace about deepfakes, how to verify them, and the value in carrying a bit of skepticism when interacting with content online. And while it will take time to see the true impact of these initiatives, Taiwan remains a democratic country with free and fair elections. For this reason, countries like the Netherlands should consider the application of similar strategies domestically.

The challenges associated with AI-driven election interference in the Netherlands are layered and there’s no one-size-fits-all solution to protecting democratic institutions and voter autonomy in the wake of deepfakes, algorithmic manipulation, and ill-informed LLMs. However, critical work is being done by consortiums like the Hybrid Election Integrity Observatory to highlight the risks to Dutch society and propose longer term solutions. As suggested by the

¹⁶ Global Taiwan Institute. “Securing US Elections from AI-Enhanced Foreign Influence: Lessons from Taiwan’s 4C Strategy.” October 30, 2024. <https://globaltaiwan.org/2024/10/securing-u-s-elections-from-ai-enhanced-foreign-influence/>.

Dutch HEIO and government, as well as foreign bodies that have seen glimmers of success in Taiwan's 4C strategy, this is an "all-hands-on-deck" situation.

Three elements distinguish the Netherlands from other less-equipped nations in responding to AI-driven interference in democratic processes: its coalition style democracy, which promotes greater diversity of political parties and resistance to choice collapse, its uniquely European approach to free speech, which recognizes the need to balance responsibility and care with legal action as needed, and finally, its ability to engage and act on complex challenges like those presented by AI-driven content. For these reasons, the Dutch are already set up to be more successful than other countries in addressing the harms caused by AI-driven content. Through a combination of interventions like real-time monitoring, increased criminal liabilities, algorithmic accountability, and more humans-in-the-loop, the Netherlands may be able to avoid the worst possible harms that result from AI-driven interference in its democratic state affairs.

Bibliography

- Car, Polona, and Beatrix Immenkamp. *Hate Speech: Comparing the US and EU Approaches*. PE 772.890. European Parliamentary Research Service (EPRS), European Parliament, 2025.
[https://www.europarl.europa.eu/RegData/etudes/BRIE/2025/772890/EPRS_BRI\(2025\)772890_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2025/772890/EPRS_BRI(2025)772890_EN.pdf).
- Carpenter-Zehe, Olivia. "AI-Driven Digital Manipulation 'tested' Dutch Election Integrity, Researchers Warn." January 16, 2026. <https://euobserver.com/197303/ai-driven-digital-manipulation-tested-dutch-election-integrity-researchers-warn/>.
- de León Williams, Ernesto, and Fabio Votta. "1 in 10 Dutch Citizens Are Likely to Ask AI for Election Advice. This Is Why They Shouldn't." *The University of Amsterdam*, October 29, 2025. <https://www.uva.nl/en/shared-content/faculteiten/en/faculteit-der-maatschappij-en-gedragswetenschappen/news/2025/10/dont-ask-ai-for-election-advice.html>.
- DutchNews.nl. "Beware the Trolls, Institute Tells Government as Elections Near." March 11, 2026. <https://www.dutchnews.nl/2026/03/beware-the-trolls-institute-tells-government-as-elections-near/>.
- Ex, Luuk, Quirine Van Eeden, Wouter Nieuwenhuizen, and Mariette Van Huijstee. *Scrollend Naar de Stembus: De Rol van Aanbevelingsalgoritmen Bij Inmenging in Verkiezingen*. Rathenau Instituut, 2026.
https://www.rathenau.nl/sites/default/files/2026-03/Scrollend_naar_de_stembus_Rathenau_Instituut.pdf.
- Global Taiwan Institute. "Securing US Elections from AI-Enhanced Foreign Influence: Lessons from Taiwan's 4C Strategy." October 30, 2024.
<https://globaltaiwan.org/2024/10/securing-u-s-elections-from-ai-enhanced-foreign-influence/>.
- HEIO Consortium. *HEIO - Dutch General Elections Oct '25 - Final Report*. 2026.
<https://www.heio.nl/wp-content/uploads/2026/01/260115-HEIO-Final-Report.pdf>.
- Institute for Strategic Dialogue, Alliance4Europe, Debunk.org, GMF Alliance for Securing Democracy, DEN Institute, & EU DisinfoLab. *Foreign Information Manipulation in the 2025 German Federal Election*. Alliance4Europe, 2025.
<https://alliance4europe.eu/foreign-information-manipulation-2025-german-federal-election>.

- Netherlands Institute for Multiparty Democracy & Instituut voor Publiek en Politiek. "The Dutch Political System in a Nutshell." 2008. <https://nimd.org/wp-content/uploads/2025/02/Dutch-Political-System.pdf>.
- NL Times. "Dutch Extreme Right Anti-Refugee Anthem Removed from Spotify, YouTube." *NL Times*, November 12, 2025. <https://nltimes.nl/2025/11/12/dutch-extreme-right-anti-refugee-anthem-removed-spotify-youtube>.
- Parodi, Alessandro, and Toby Sterling. "Dutch Watchdog Warns Voters against Using AI Chatbots Ahead of Election." *Reuters*, October 21, 2025. <https://www.reuters.com/technology/dutch-watchdog-warns-voters-against-using-ai-chatbots-ahead-election-2025-10-21/>.
- Somers, Meredith. "Deepfakes, Explained." *MIT Sloan School of Management: Ideas Made To Matter*, July 21, 2020. <https://mitsloan.mit.edu/ideas-made-to-matter/deepfakes-explained>.
- Tudor, Grant. "Polarization, Extremism, and Proportional Representation." September 11, 2023. <https://protectdemocracy.org/work/polarization-extremism-and-proportional-representation/>.
- World Economic Forum. *The Global Risks Report 2026*. 21st ed. Geneva: World Economic Forum, January 14, 2026. https://reports.weforum.org/docs/WEF_Global_Risks_Report_2026.pdf.