FINAL REPORT

HE
IO

WWW.HEIO.NL

HYBRID ELECTION INTEGRITY OBSERVATORY

DUTCH PARLIAMENTARY ELECTIONS 2025 REPORT

OBSERVATORY PARTNERS

POST-X SOCIETY

AI FORENSICS

TROLLRENSICS

UNIVERSITY OF AMSTERDAM

JUSTICE FOR PROSPERITY

WWW.HEIO.NL

OVERVIEW

About the Hybrid Election Integrity Observatory

The Hybrid Election Integrity Observatory (HEIO) is a consortium project that protects election integrity by monitoring social media platforms for interference operations. During the Dutch parliamentary elections campaign period leading up to October 29th 2025, HEIO combined cross-platform monitoring capabilities with field investigations and journalistic collaborations to detect, analyze, and counter threats to electoral integrity.

COLOPHON

Title:	HEIO - Dutch General Elections Oct ‘25 - Final Report
Subtitle:	Monitoring the Dutch Parliamentary Elections, October 29, 2025
Published:	January 2026
Authors:	HEIO Consortium - Post-X Society, AI Forensics, Trollrensics, University of Amsterdam, Justice for Prosperity

Cite as

HEIO Consortium (2026). Hybrid Election Integrity Observatory: Final Report - Monitoring the Dutch Parliamentary Elections, October 29th 2025. The Hague: Post-X Society

License

EXECUTIVE SUMMARY

The Dutch parliamentary elections of October 29, 2025, took place amid heightened concerns about digital interference, particularly following reports from other European election monitors and Romania's cancellation of their presidential election due to foreign manipulation on social media. The Hybrid Election Integrity Observatory (HEIO) was established as an immediate response to these evolving threats, bringing together five specialized organizations to monitor and analyze election integrity across social media platforms.

Our findings confirm that while the Dutch elections remained fundamentally free and fair, they were conducted under significant digital pressure. The election landscape was characterized by a surge in AI-generated content, coordinated manipulation campaigns, platform moderation failures, and the emergence of new ways for spreading disinformation and hate speech.

The HEIO consortium documented multiple instances of coordinated inauthentic behavior. On three platforms our consortium partner Trollrensics identified large networks. Around 23,000 accounts from Vietnam where detected and analysed, which massively placed likes on the Facebook page of GL-PVDA leader Frans Timmermans for about a month. On X thousands of accounts from Nigeria, Ghana and Ivory Coast have been found retweeting polarizing content and content of far right political parties PVV and FvD. On Youtube a troll operation has been identified, but the analysis is only in its start phase and there are only limited results.

Next, we documented an enormous growth of generative AI content. An AI-generated protest song "Wij zeggen nee, nee, nee tegen een AZC" reached #2 on Spotify's Netherlands Top 50, spawning over thousands TikTok videos in a matter of weeks. AI-generated images became progressively more extreme throughout the campaign period. Some crossing into potentially illegal territory depicting violence against politicians and hateful content against minorities.

The most egregious example was a coordinated Facebook page operation that became the most popular political page in the Netherlands, sometimes reaching over one million daily views, and totaling 74 million views between June and October 2025.

"This election was also the least transparent in recent Dutch history"

Platform moderation proved inadequate: not a single piece of content reported through official platform reporting mechanisms was removed, even when content clearly violated platform terms of service. Only when content was exposed through media coverage did platforms respond. This indicates that public embarrassment, rather than user safety, drives enforcement by the platforms.

Livestreams emerged as particularly problematic spaces where death threats, antisemitism, and explicit racism flourished without intervention, especially on TikTok. The ephemeral nature of live content and lack of effective monitoring created accountability-free zones for hate speech and incitement. Meanwhile, very new accounts (less than three days old) with minimal content appeared prominently in algorithmic feeds, suggesting systematic manipulation of recommendation systems.

A surge in online narratives regarding election fraud took place, as we anticipated. From election day on forward >2.000 messages related to this topic were posted on X. When Geert Wilders choose to amplify baseless claims, all media responded quickly to refute these.

This election was also the least transparent in recent Dutch history regarding political advertising. The EU regulation requiring political ad transparency resulted in major platforms from Meta and Google banned political ads all together. The platforms that did allow political ads poorly implemented the policies, with ad libraries remaining incomplete and incomparable across platforms. X, TikTok and Snapchat ad databases were particularly inadequate. The voluntary self-reporting system for digital political advertising in other spaces proved completely ineffective, providing no meaningful oversight of campaign spending or messaging.

Our interactions with Dutch and European authorities revealed systemic weaknesses in the current regulatory framework. The Dutch DSA coordinator (ACM) is powerless without formal complaints, and the EC Rapid Response System comes with contractual secrecy, rendering us incapable of providing transparency about outcomes. Research data access requests under the DSA came too late to be useful during the election period, regardless of our timely requests.

Most fundamentally, this project exposed an untenable situation: the monitoring of systemic risks to democratic processes is conducted by our non-governmental organizations with very limited resources, while platforms get away with soft promises without meaningful accountability mechanisms.

For as far as we can tell, the elections were free and fair, but they were also under threat. The digital information ecosystem surrounding Dutch democracy is fragile, poorly regulated, and increasingly vulnerable to manipulation. Without structural changes to platform accountability, funding for adequate and rapid monitoring infrastructure, and enforcement mechanisms that work in real-time rather than retrospectively, future elections face growing risks. The techniques and networks identified during this monitoring period remain active and will most likely target future municipal, provincial, and European elections.

Based on our observations we formulated eleven recommendations that could immediately improve election integrity, focussing on accountability, transparency and regulatory reform.

MANAGEMENT SAMENVATTING

De Tweede Kamerverkiezingen van 29 oktober 2025 volgden kort na berichtgeving over digitale inmenging in andere verkiezingen in Europa, met name de Roemeense presidentsverkiezingen. Het Hybrid Election Integrity Observatory (HEIO) werd opgericht als reactie op deze ontwikkelingen. Vijf gespecialiseerde organisaties monitorden de integriteit van de verkiezingscampagne op sociale mediaplatforms en rapporteren hier hun observaties.

De Nederlandse verkiezingen verliepen vrij en eerlijk, maar stonden online onder aanzienlijke druk. Online campagnes maakten op grote schaal gebruik van AI-gegenereerde content. Verder zagen we gecoördineerde manipulatiecampagnes, het falen van platformmoderatie, en het benutten van nieuwe technieken om desinformatie en haatzaaiende uitingen te verspreiden.

Onze consortiumpartner Trollrensics documenteerde meerdere gevallen van gecoördineerd inauthentiek gedrag. Ruim 23.000 accounts ingekochte bij een Vietnamese trollenfabriek die werden ingezet om content van Frans Timmermans op Facebook te liken. Op X spoorden we duizenden acounts op die vanuit Nigeria, Ghana en Ivoorkust polariserende content retweeten, en berichten van PVV en FvD. Op YouTube is het onderzoek nog lopende.

Opvallend was verder de enorme toename in het gebuik van generatieve AI-content. Een AI-gegenereerd protestlied "Wij zeggen nee, nee, nee tegen een AZC" bereikte de #2-positie in de Spotify Top 50 van Nederland en duizenden TikTok-video's gebruikten het als achtergrond binnen enkele weken. AI-gegenereerde beelden werden gedurende de campagneperiode steeds extremer. Sommige berichten zijn mogelijk illegaal vanwege het afbeelden van geweld tegen politici en haat tegen minderheden.

Het meest schokkende voorbeeld was een gecoördineerde Facebook-campagne die de populairste politieke pagina van Nederland omvatte. Deze pagina trok soms meer dan een miljoen views per dag. In totaal werden deze beelden 74 miljoen keer bekeken tussen juni en oktober 2025.

> "Deze verkiezing waren de minst transparante in de recente Nederlandse geschiedenis "

De rapportagefuncties van platformen bleken niet te werken. Niet één bericht dat wij meldden via de officiële meldmechanismen van platforms werd verwijderd, zelfs wanneer content duidelijk de gebruikersvoorwaarden schond. Pas toen we gewelddadige AI-gegenereerde content via media-aandacht onthulden, reageerden platforms door te modereren. Dit wijst erop dat het beperken van reputatieschade voor hen belangrijker is dan de veiligheid van gebruikers.

Livestreams bleken bijzonder problematische online omgevingen te zijn. Op TikTok observeerden we doodsbedreigingen, antisemitisme en expliciet racisme zonder interventie. De vluchtige aard van live-content en gebrek aan effectieve monitoring creëerden een soort vrije zones voor haatzaaien en aanzetten tot geweld. Verder zagen we zeer nieuwe accounts (soms minder dan drie dagen oud) met minimale content prominent in algoritmische feeds verschijnen, wat systematische manipulatie van aanbevelingssystemen suggereert.

Deze verkiezing waren ook de minst transparante in de recente Nederlandse geschiedenis wat betreft politieke advertenties. Vanwege de invoering van EU-regelgeving voor transparantie van politieke advertenties besloten Meta en Google politieke advertenties te verbieden. Andere platformen hebben regelgeving slecht geïmplementeerd. Advertentiebibliotheken waren incompleet en onvergelijkbaar tussen platforms. De databases van X, TikTok en Snapchat waren bijzonder intransparant en ongebruiksvriendelijk. Verder schoot het vrijwillige zelfrapportagesysteem voor politieke advertenties in andere digitale omgevingen tekort. Het boodt geen zinnig inzicht op campagneuitgaven of boodschappen.

Onze interacties met Nederlandse en Europese autoriteiten brengen ook systemische zwaktes in de huidige regelgeving aan het licht. De Nederlandse DSA-coördinator (ACM) kan niet handhaven zonder formele klachten. Onze medewerking aan het Rapid Response System van de Europese Commissie gaat gepaard met een geheimhoudingsclausule. Hierdoor kunnen wij geen inzicht bieden in de resultaten van onze meldingen. Verder werd ons verzoek om onderzoeksdatatoegang te verkrijgen onder de DSA te laat ingewilligd om van nut te zijn.

Tot slot brengt ons observatorium een onhoudbare situatie aan het licht. De monitoring van systemische risico's voor onze democratische processen is op dit moment afhankelijk van door maatschappelijke organisaties met beperkte middelen. De platforms kunnen vrijuit zachte beloftes doen, gratis gebruik maken van het werk dat wij verzetten, en hoeven nauwelijks verantwoording af te leggen.

Voor zover wij konden oordelen, verliepen de verkiezingen vrij en eerlijk, maar dat stond wel onder druk. De digitale kant van de verkiezingscampagne is kwetsbaar gebleken, slecht gereguleerd en vatbaarder voor manipulatie. Structurele veranderingen in verantwoordelijkheden van platformen, structurele financiering voor tijdige en adequate monitoring en handhaving die real-time optreed is nodig. Anders lopen toekomstige verkiezingen in toenemende mate risico's. De technieken en netwerken die wij hebben geobserveerd, zullen actief blijven. Niets weerhoudt actoren ervan om te proberen toekomstige gemeentelijke, provinciale, landelijke en Europese verkiezingen te verstoren.

Op basis van onze ervaringen hebben we elf aanbevelingen geformuleerd die direct kunnen bijdragen aan het weerbaar maken van verkiezingscampagnes, gericht op verantwoordelijkheid, transparantie en herziening van wet- en regelgeving.

Meet the observatory partners

Post-X Society
PXS

Coordination and

TROLLRENSICS

Detecting and analyzing coordinated inauthentic behavior and information operations

AI FORENSICS
AIF

Nobis urbi autendenihil iducias obis autendenihil as iducias m quidipsa quam recus. Quas dolentiur iducias?

Justice for Prosperity
JfP

Nobis urbi autendenihil iducias obis autendenihil as iducias m quidipsa quam recus. Quas dolen-

University of Amsterdam
UvA

Nobis urbi autendenihil iducias obis autendenihil as iducias m quidipsa quam recus. Quas dolen-

Advisory Board

The HEIO project benefited from guidance by an international advisory board of experts in social media, content creation, and disinformation:

- Marcus Bösch - University of Münster, Newsletter Understanding TikTok
- Tom Divon - Content & Creators Researcher, Dept. of Communication & Journalism, The Hebrew University of Jerusalem, co-founder The Content Creator Scholars Network
- Lea Frühwirth - Senior researcher on Disinformation CeMAS (Center für Monitoring, Analyse und Strategie)
- Esther Hammelburg - Amsterdam University of Applied Sciences

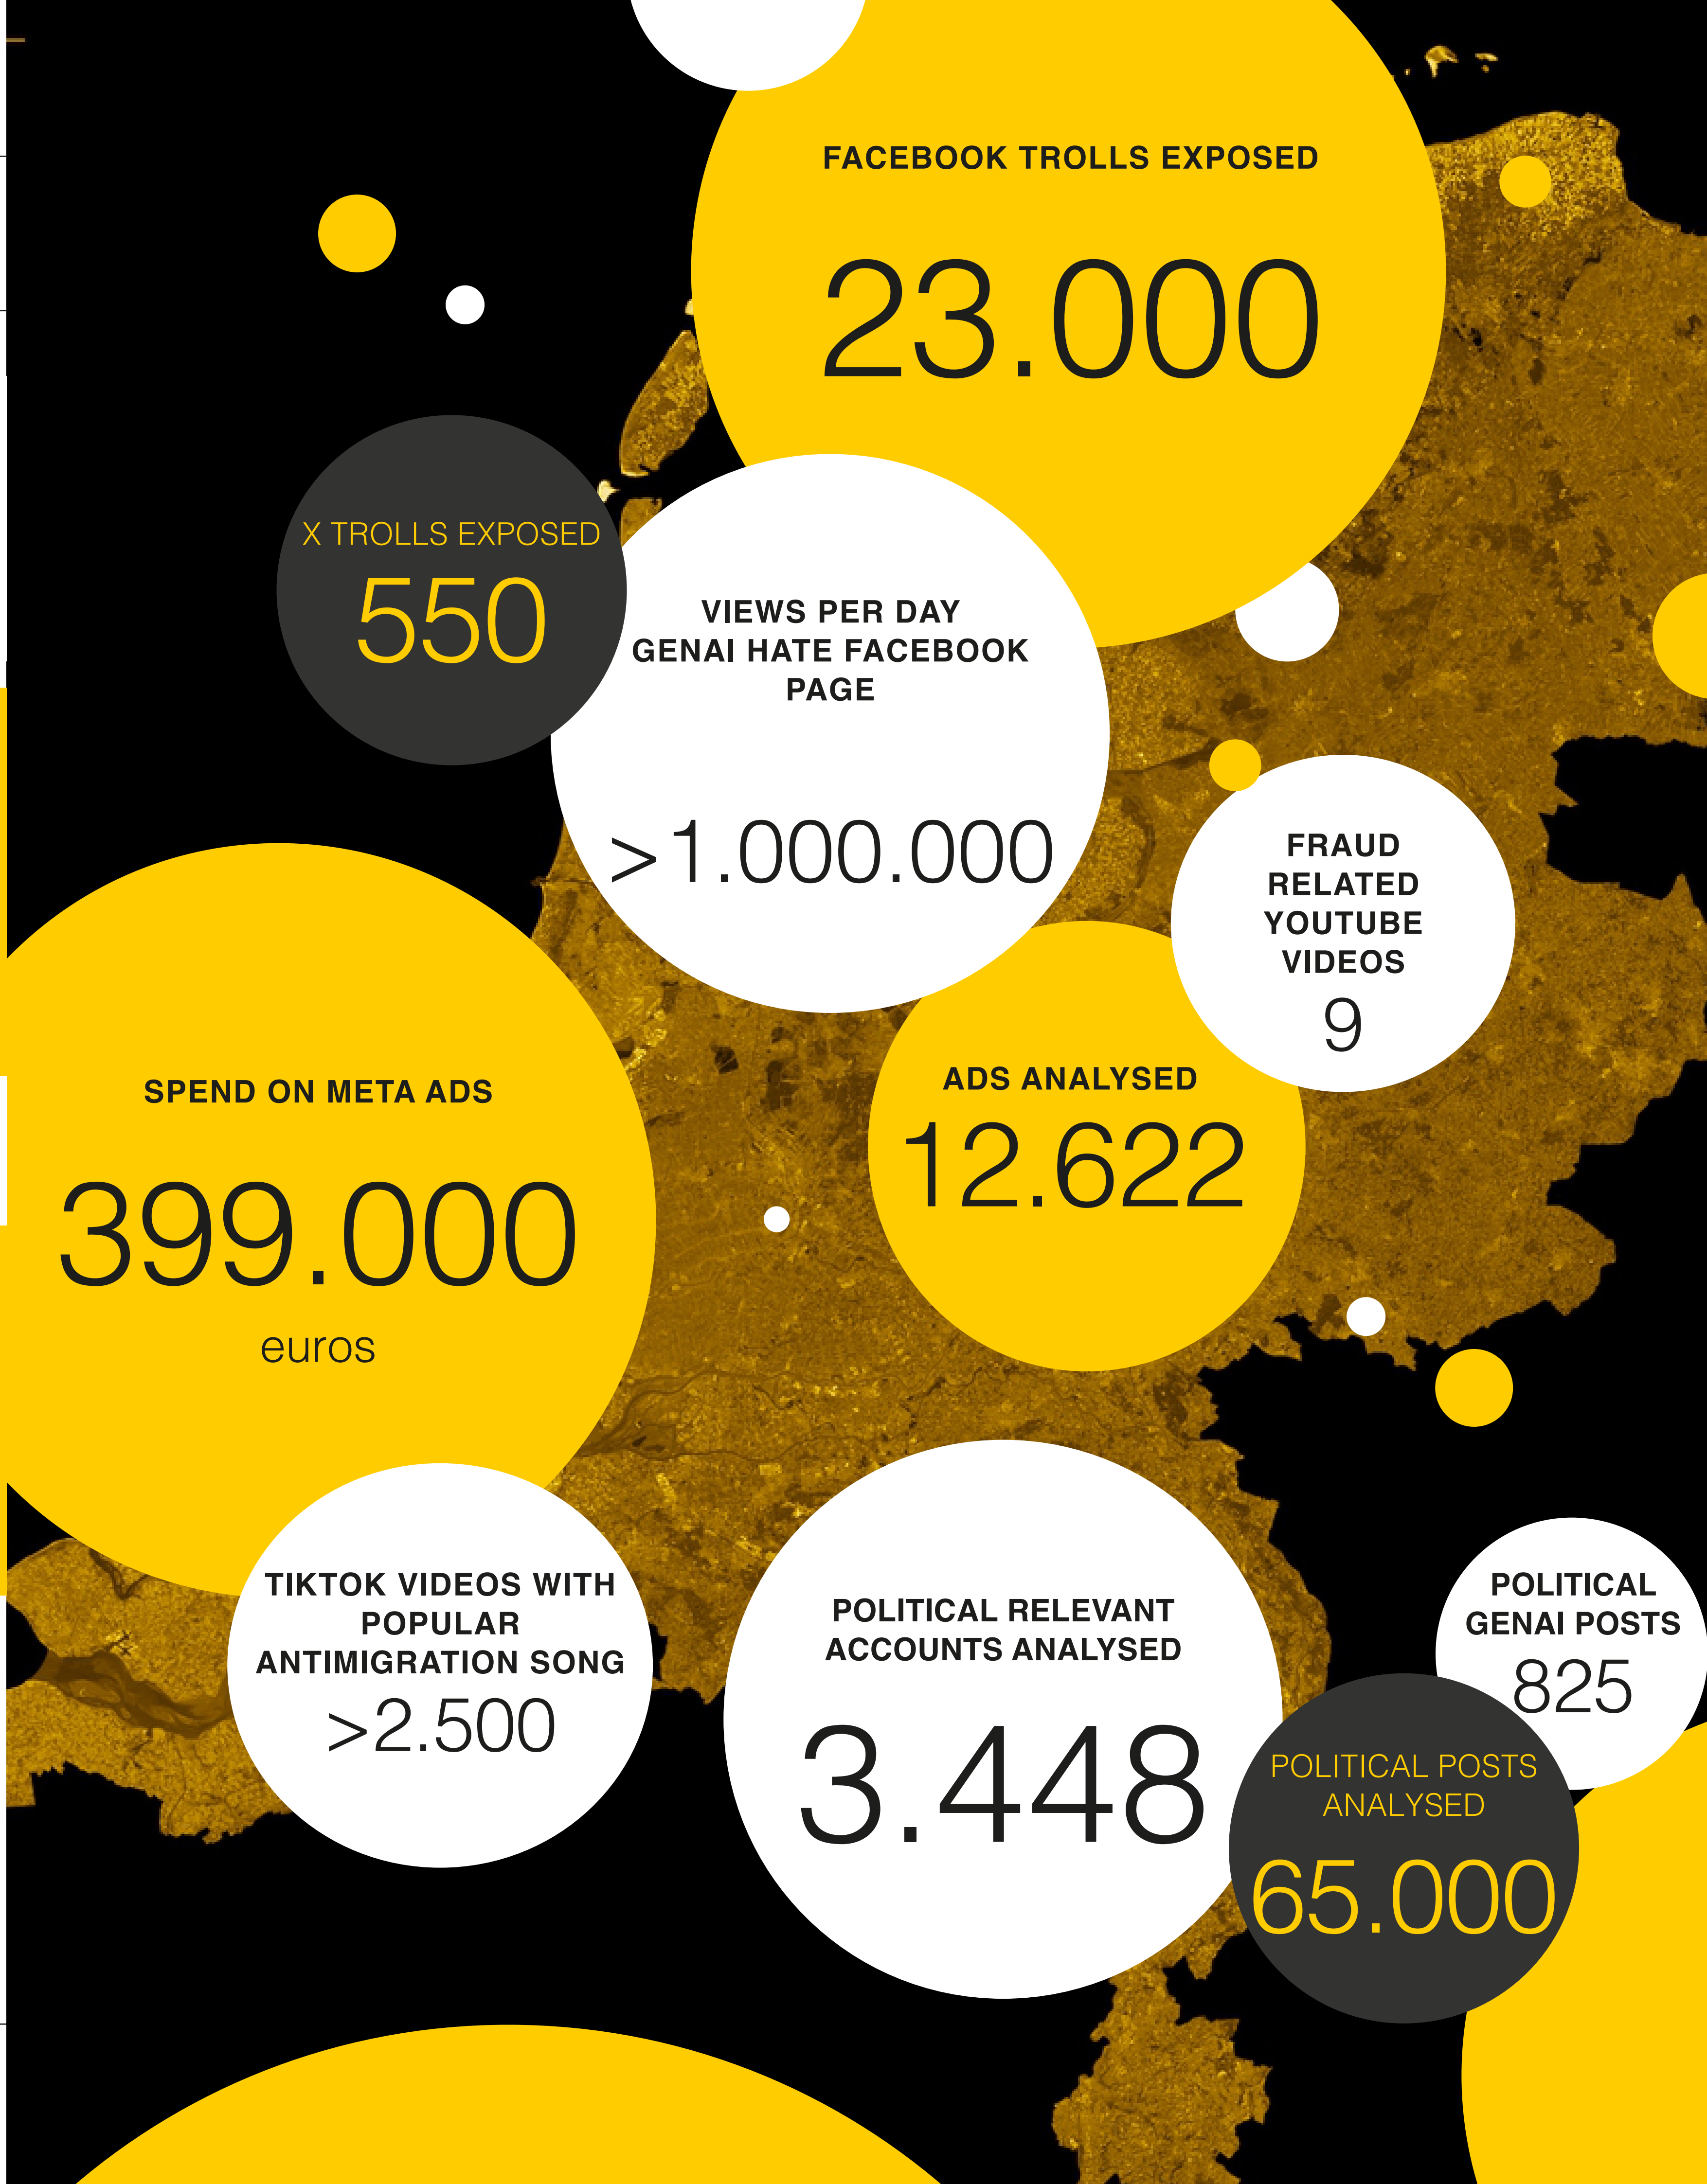

FACEBOOK TROLLS EXPOSED
23.000

X TROLLS EXPOSED
550

VIEWS PER DAY GENAI HATE FACEBOOK PAGE
>1.000.000

FRAUD RELATED YOUTUBE VIDEOS
9

SPEND ON META ADS
399.000
euros

ADS ANALYSED
12.622

TIKTOK VIDEOS WITH POPULAR ANTIMIGRATION SONG
>2.500

POLITICAL RELEVANT ACCOUNTS ANALYSED
3.448

POLITICAL GENAI POSTS
825

POLITICAL POSTS ANALYSED
65.000

1

INTRODUCTION

INTRODUCTION

The Hybrid Election Integrity Observatory (HEIO) was established as a rapid response to escalating threats to election integrity across the globe. Our consortium of researchers and civil society organizations tracked digital threats to election integrity throughout the campaign period of the Dutch general elections in October 2025.

Our primary objectives were:

Detection
Identifying foreign and domestic interference operations targeting the Dutch elections across multiple social media platforms

Analysis
Document patterns of coordinated inauthentic behavior, disinformation campaigns, platform moderation failures, and algorithmic manipulation

Rapid response
Alert relevant authorities, affected parties, and the public about critical threats in time to enable protective action

Documentation
Create a comprehensive record of the digital information ecosystem during the 2025 Dutch elections for policy development and as a reference point for future research

Accountability
Hold platforms and actors responsible for violations of terms of service, Dutch law, and European regulations

HEIO brought together five specialized organizations with complementary expertise: Post-X Society (project coordination, TikTok monitoring, livestream analysis), AI Forensics (algorithmic analysis, ad library monitoring), Trollrensics (coordinated inauthentic behavior detection, network analysis), University of Amsterdam (GenAI content analysis, political advertising), and Justice for Prosperity (field investigations, OSINT research). This enabled us to monitor cross-platform with diverse methodologies, and rapidly cross-validate findings.

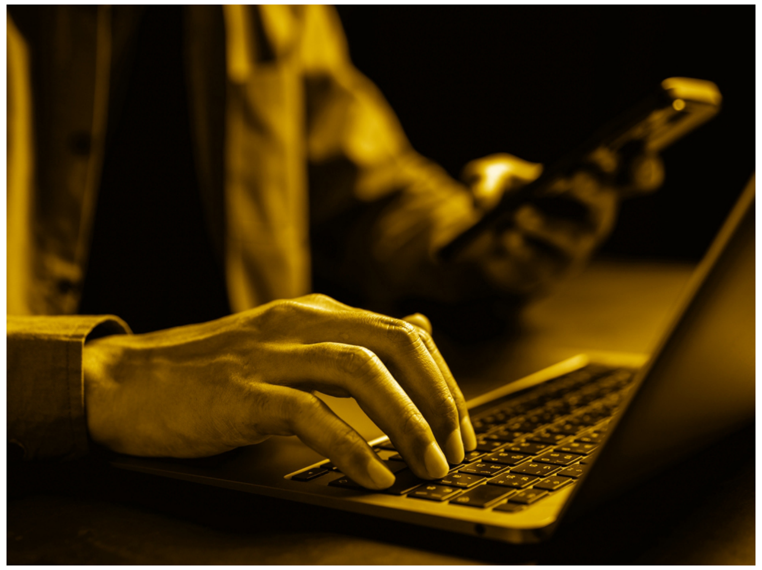

What are observations?

Before reading our report it is essential to understand the nature of our monitoring approach. We did not conduct a complete systematic analysis of all political content across all platforms during the election period. Such an undertaking would require resources, access, and time beyond what was available for this project.

Instead, we focused on identifying violations, outliers, anomalies, and patterns indicative of inauthentic behavior or malicious activity. Our monitoring was deliberately targeted at:

Content that violated platform terms of service

disinformation, hate speech, violent threats, harassment, impersonation

Violations of Dutch law and EU regulations

illegal content, DSA non-compliance, advertising transparency failures

Statistical anomalies

unusual growth patterns, coordinated timing, artificial amplification

Coordinated inauthentic behavior

networks of fake accounts, purchased engagement, cross-platform manipulation

Algorithmic irregularities

suspicious recommendations, search results or visibility

About this report

This final report documents HEIO's findings from the Dutch parliamentary election period, including the weeks immediately preceding and following October 29, 2025. We report on observations from multiple monitoring streams, interactions with authorities and platforms, and analysis of emerging threats to electoral integrity.

The structure is as follows. The main findings are summarized at the beginning, followed by a short description of the methodology and a more comprehensive descriptions of observations, interactions with authorities, conclusions, and recommendations for protecting future elections.

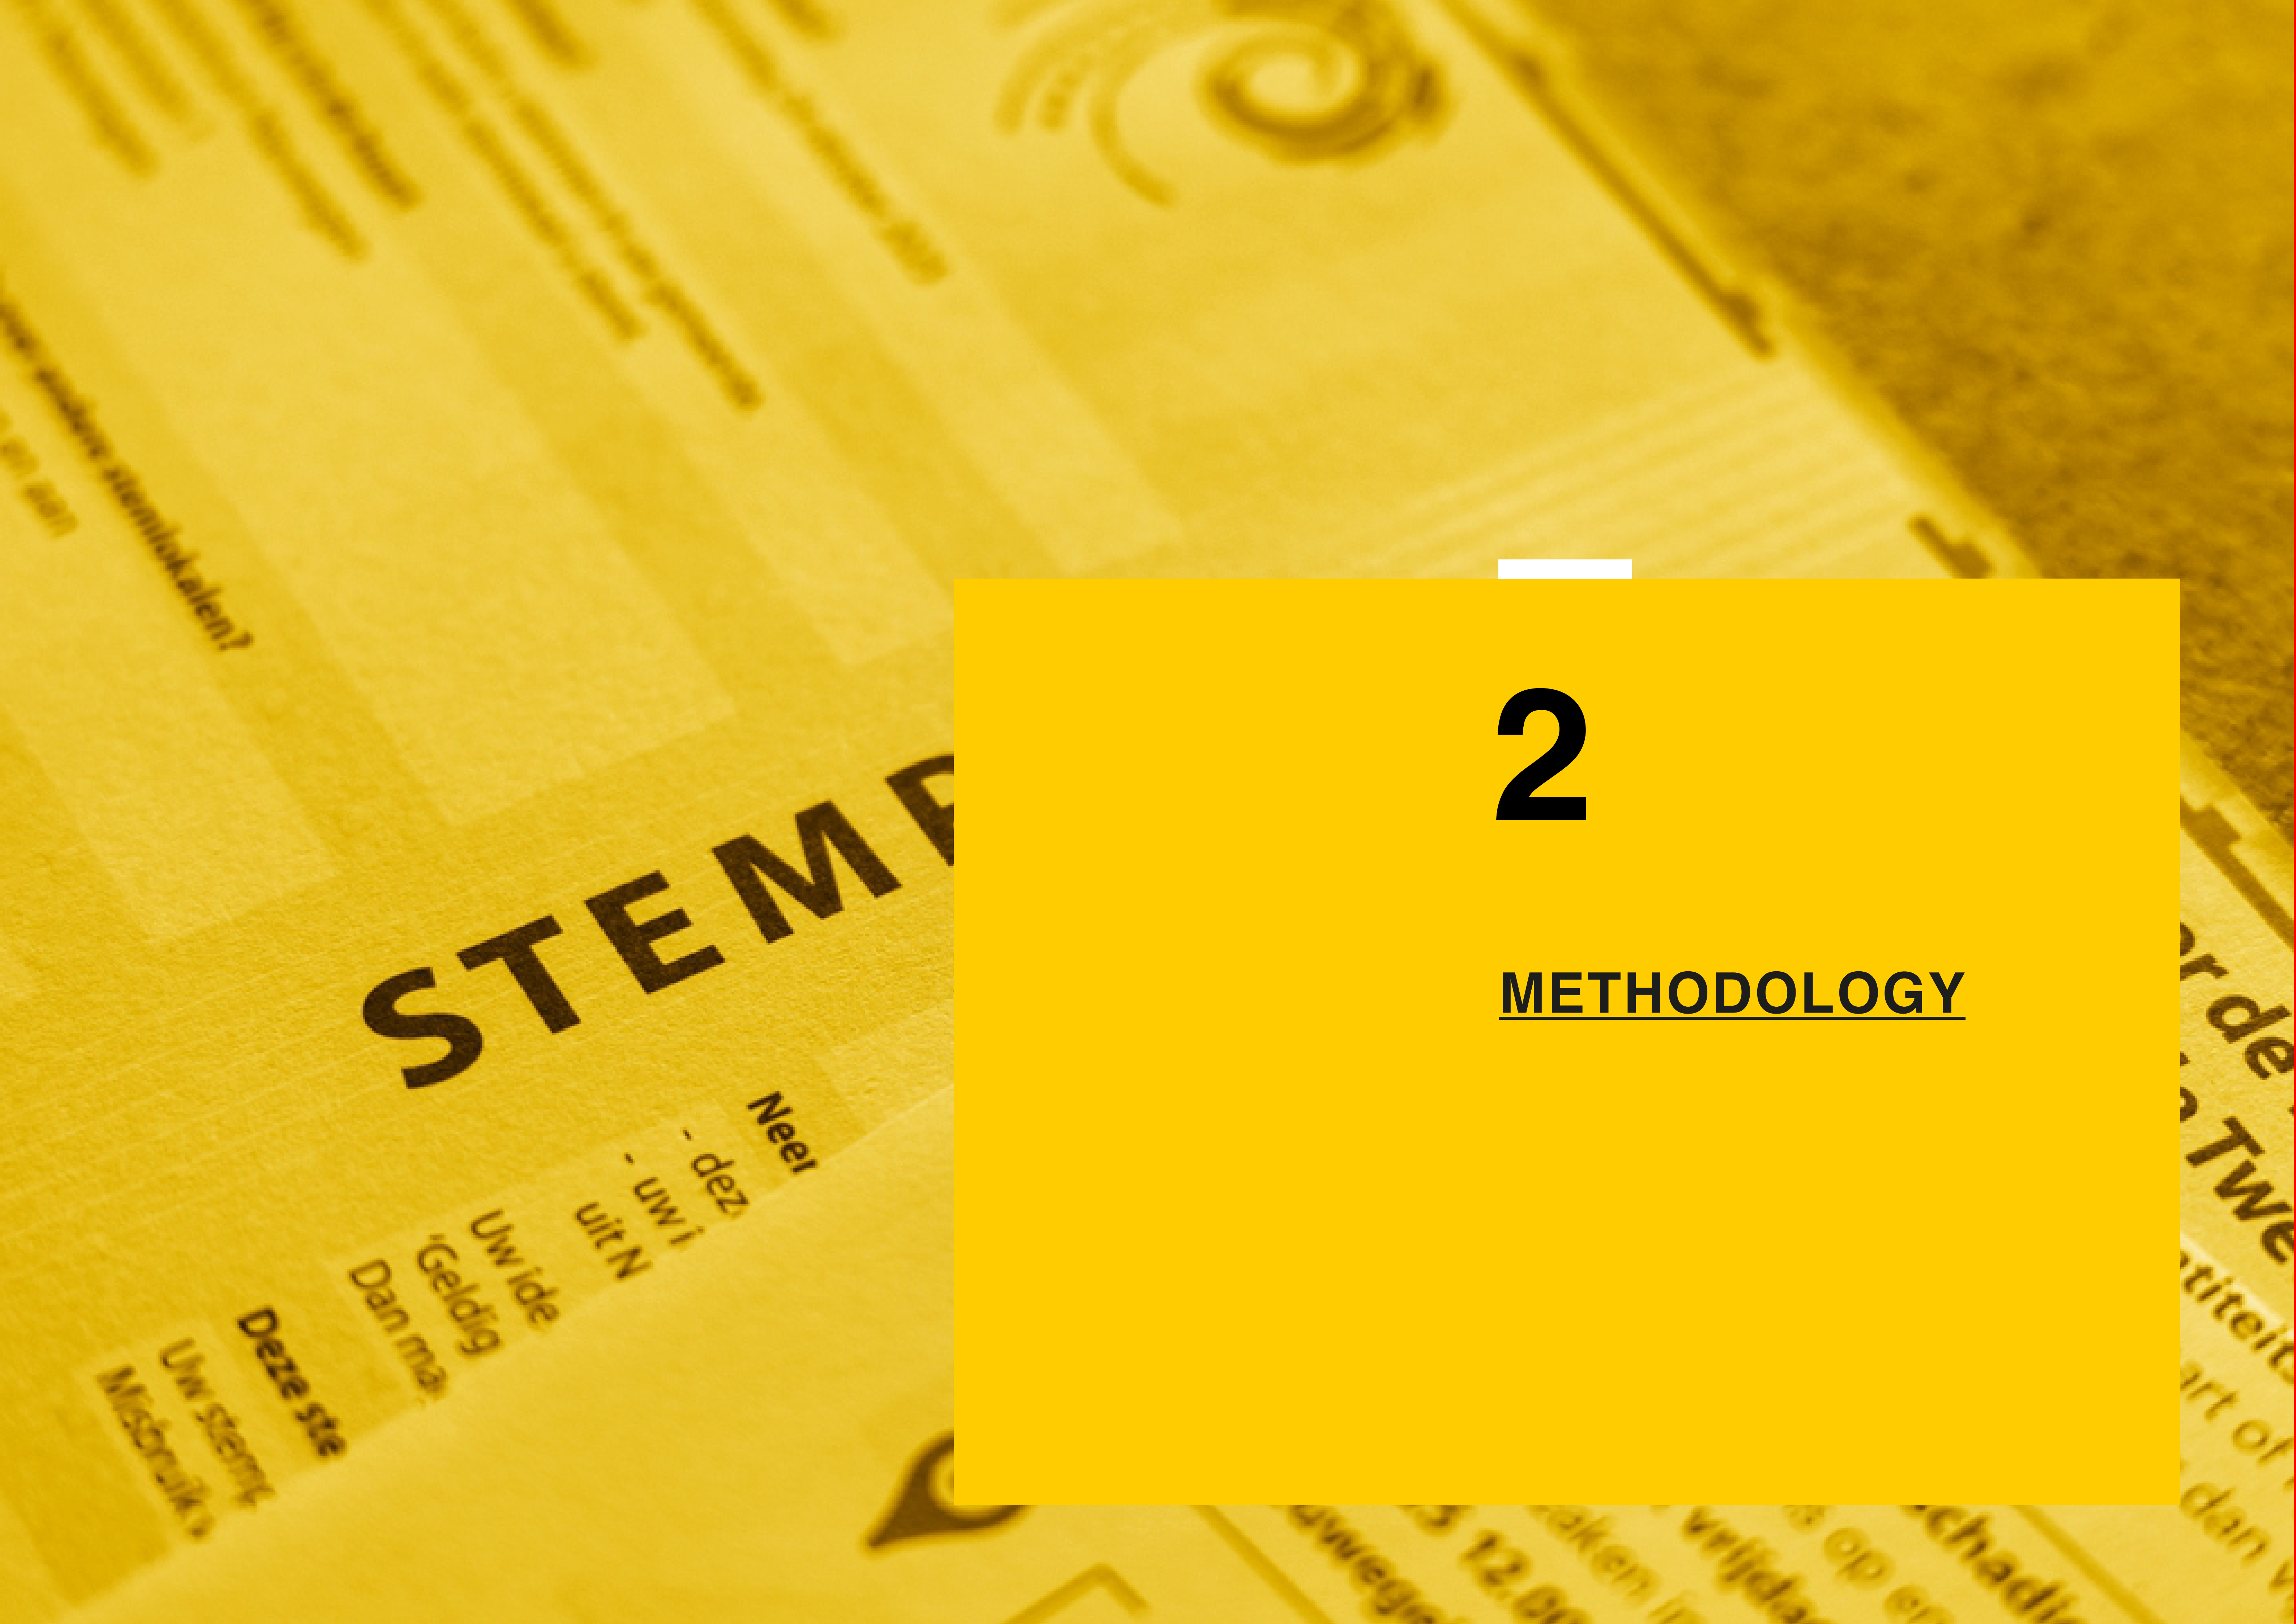

2

METHODOLOGY

Our Methods

HEIO employed a multi-platform, multi-method approach combining automated monitoring tools, manual analysis, field investigations, and journalistic collaboration.

» PLATFORM MONITORING «

TROLLRENSICS PLATFORM

We monitored X, Telegram, and TikTok for coordination patterns and influence operations using several applications, including the Trollrensics platform. We tracked keywords, hashtags, and accounts to detect coordinated inauthentic behavior and cross-platform campaigns. These included generic election related terms, trending topics and accounts of political parties and politicians.

To enable journalist to conduct research, Trollrensics provided software licenses to NRC and RTL. Two NRC and two RTL journalists were trained in using the software. Trollrensics started analyzing data from about the beginning of September 2025.

» GENAI MONITORING «

CampAIgn Tracker

We developed a dedicated visual AI tracking infrastructure, building on our earlier work for the German elections (campaigntracker.de), in collaboration with Simon Kruschinski (Senior Researcher, GESIS - Leibniz Institute for the Social Sciences). For the Dutch parliamentary elections, we deployed a GenAI dashboard that continuously collected and analysed public content from Facebook, Instagram, TikTok and X for parties, candidates, political commentators, meme pages, and politically relevant influencers. We assembled a list of 3,448 accounts at national, provincial, and local levels.

All public posts from these accounts within the election period were ingested into our detection pipeline, together with basic metadata (account, timestamp, platform, engagement) including the media used (images and videos). For the Dutch elections, this resulted in over 65.000 posts between 17 September 2025 and 29 October 2025. Accounts were classified and validated by our research assistants and volunteers from Who Targets Me, who contributed their existing mappings of party and advertiser accounts from previous election cycles.

CampAIn Tracker GenAI detection

To identify AI-generated and AI-manipulated visuals, UvA used a two-stage detection process:

1. Automated pre-screening

All collected images and video stills were scored using the SightEngine AI-detection model. Any item with a model probability above 0.1 of being AI-generated was flagged for human review.

2. Human validation and manual coding

Trained research assistants manually checked all flagged items to confirm whether AI was actually used. For all confirmed AI posts, coders then applied a detailed coding scheme capturing:

- whether AI use was explicitly labelled or disclosed
- which actors were depicted (e.g. specific politicians, parties, social groups)
- the main theme (e.g. migration, crime, housing)
- the actors mentioned in the full post (caption, linked text)
- the use of negativity (attacks, delegitimisation) and acclaims (self-praise, success claims).

Using this process we identified 852 posts containing AI imagery. These data feed into public-facing visualisations on campaigntracker.nl that allow citizens to explore where and how GenAI was used in the campaign, and which actors or themes it was attached to.

» AD MONITORING «

We conducted systematic monitoring of political ad repositories across Meta, TikTok, Snapchat, X, and voluntary self-reporting databases. UvA, JfP and AIF used complementary methods.

UVA AD MONITORING

To complement the content analysis, we built dashboards that track both political ad spending and targeting practices. For spending, we scraped and harmonised data from three Dutch transparency sources, politiekereclame.nl, Ster, and DPG Media, using the open-source R package reclamer. The spending dashboard is publicly accessible.

For each campaign, we standardised advertiser names, campaign periods, channels, and reported amounts (approximating DPG's spending brackets with midpoints) to produce comparable estimates of total spend by party, actor type, media outlet, and channel. These figures are informative but necessarily imperfect, as they rely on self-reported and sometimes incomplete data.

UvA collaborated with Who Targets Me to analyse Meta's Ad Library "Audience" data for 999 political advertisers across multiple elections. Using the R package metatargetr, we regularly retrieved rolling 7-, 30-, and 90-day targeting windows and aggregated them to show what share of each advertiser's budget relied on broad location-only targeting, detailed targeting, custom audiences, or lookalike audiences. Because Ad Library data are known to be incomplete and delayed, all targeting indicators are presented as lower-bound estimates. The targeting dashboard is accessible here. Together, the spending and targeting dashboards help us see not only what political content was promoted, but also how money and data were used to reach specific audiences.

JFP AD MONITORING

Justice for Prosperity used the advertising repositories of Meta, TikTok, Snapchat, X to conduct daily searches on a predefined set of election-related keywords to conduct a systematic monitoring. These searches allowed us to determine whether political or politically targeted advertisements were active on each day. When an advertisement seemed to be particularly noteworthy, for example when containing misleading or false information, we carried out further investigation using OSINT techniques. This method allowed JfP to track polarising actors who use advertisements and take advantage of the platform's DSA non-compliance to further spread misinfomation. This monitoring allowed us to distinguish three separate types of advertisements: genuine political ads, commercial ads which hijacked the political terminology and lastly, false or scam advertisements.

AIF AD MONITORING

AIF queried the Meta Ad Library API with a set of 46 keywords, containing the names of the lead candidates in the elections, as well as their corresponding political parties and abbreviations (for example: Democrats 66, D66 and Rob Jetten). This search returned 12.622 advertisements, which were refined by limiting the dataset to only ads posted in The Netherlands and posted in the period from October 16th to 30th of 2025. Since our collection was done after Meta's policy went into effect, 653 ads were labeled as "removed by Meta for violating standards," by Meta. Excluding those already flagged by Meta, we filtered the ads that contained the aforementioned party and candidate keywords in the accounts name, and found 237 advertisements, which included non-exact keyword matches. A researcher manually reviewed the 237 advertisements to find the relevant candidate and political party accounts.

» YOUTUBE MONITORING «

MONITORING ELECTION FRAUD NARRATIVES ON YOUTUBE

AIF assessed the presence of election fraud on Youtube by using a list of 23 Dutch keywords related to election fraud, such as "Verkiezingen fraude" or "Stemmen fraude", to perform daily queries on the platform. This list was informed by Dutch researchers and journalists familiar with election fraud narratives in The Netherlands. Between October 21st and November 1st 2025 we repeated each query every hour, collecting the top 60 results for each query, totaling 1.171.022 collected search results, which contained 15.199 unique videos.

Since we were interested in what a Dutch person would see, we built a custom scraper leveraging Dutch residential IP proxies to perform the queries on the platform. This allowed us to see videos presented and ordered in the same way a user of the web platform would see.

We filtered the dataset for videos published after 20 October 2025 and in either Dutch or English, under the assumption that those would be the most relevant. We aimed to create two comparative datasets, one with videos that speak generally about the Dutch elections, and another that specifically focuses on fraud allegations in the Dutch elections with the intention to understand how YouTube amplifies fraud allegation videos compared to general election videos.

For the first dataset, that of general Dutch election videos, we filtered for videos that mentioned "Elections", "Voting", "Polling station", "Ballot box", "Tweede Kamerverkiezingen", "Verkiezingen", "Stemmen", "Stembureau", "Stembus", or the relevant political party names in their captions, which resulted in a dataset of 1.000 videos. From there, we ran the dataset through an LLM - llama-3.3-70b-instruct from Meta provided by OpenRouter - with the prompt to flag any videos that spoke specifically about the Dutch elections, leaving us with a dataset of 486 videos.

The second dataset, which focused on videos with fraud allegations, followed similar steps to the aforementioned dataset. Rather than filtering for election-related keywords, we filtered for words related to election fraud: "fraud", "fraude", "stolen", "rigged", "manipulated", "falsified", "gestolen", "gemanipuleerd", "bedrog", "diefstal". There were a total of 468 videos in this dataset, which we reduced to nine by using the same LLM to classify which videos discussed fraud in the context of the 2025 Dutch elections, and then we involved two researchers who manually reviewed the content.

» OSINT Tools «

OSINT TOOLS

Partners employed professional OSINT tools including but not limited to Meltwater, reverse image search, Security Trails, BuiltWith, WHOIS, Hunter.io, Internet Archive, and custom-developed tools. Persona accounts were used to access content and platform features otherwise unavailable to researchers.

» TIKTOK MONITORING «

TIKTOK LIVE STREAM MONITORING

PXS continuously monitored TikTok livestreams, automatically analyzing comments, ttracking patterns in gifting, and detecting extreme speech. The monitoring was conducted using a custom build tool: TikTokLive_Monitor. The transcription wasconducted with Whisper, and analysis with Detoxify and Claude Sonnet 4.5 with custom prompts.

PERSONA ACCOUNTS

In order to observe what users from a range of political spectra get recommended in their feeds, such as the For You Page on TikTok, we created persona accounts. These accounts were trained to recommend content with a particular political leaning, by manually scrolling and watching content that align with that orientation.

DIGITAL SERVICES ACT RESEARCHER DATA ACCESS

To further investigate the online content that could potentially spread disinformation or was created for other harmful purposes, we have made use of researcher data access provided under the EU Digital Services Act (DSA).

META ACCESS REQUEST DATE
Aug 28th

ELECTION DATE
Oct 28th

GRANTED
Nov 5th

DSA researcher access would have allowed us to assess the usability and reliability of this relatively new data access mechanism.

For some platforms the DSA research data access process proved too slow to be useful during the election period. The best example in this case is Meta, which granted access only after elections were already concluded.

To be more specific, we have requested access under the DSA to conduct research on the ad repository on the 28th of August, so we could monitor the repository in the period leading up to the elections. However, we were granted access on the 5th of November, almost a week after the elections already took place and more than two months after the access request was submitted. The gap between the request and the granted access hindered our ability to conduct timely research.

Alternatively, other platforms, particularly X/Twitter and Snapchat simply disallowed our request for access, even though election integrity is a systemic risk according to the DSA.

Beyond non-compliance from the platforms, another reason why the DSA researcher access is not as streamlined or even usable as it ought to be for the instrument to achieve its aims is that researchers get locked into platform systems that prevent team collaboration. This is because safe rooms conditions are applied for analyzing public data. These restrictions create problems for scientific reproductability and prevent real-time monitoring during critical periods.

3

OBSERVATIONS

Observations

In this chapter we share observations made using our previously described methods

HEIO OBSERVATION TYPOLOGY

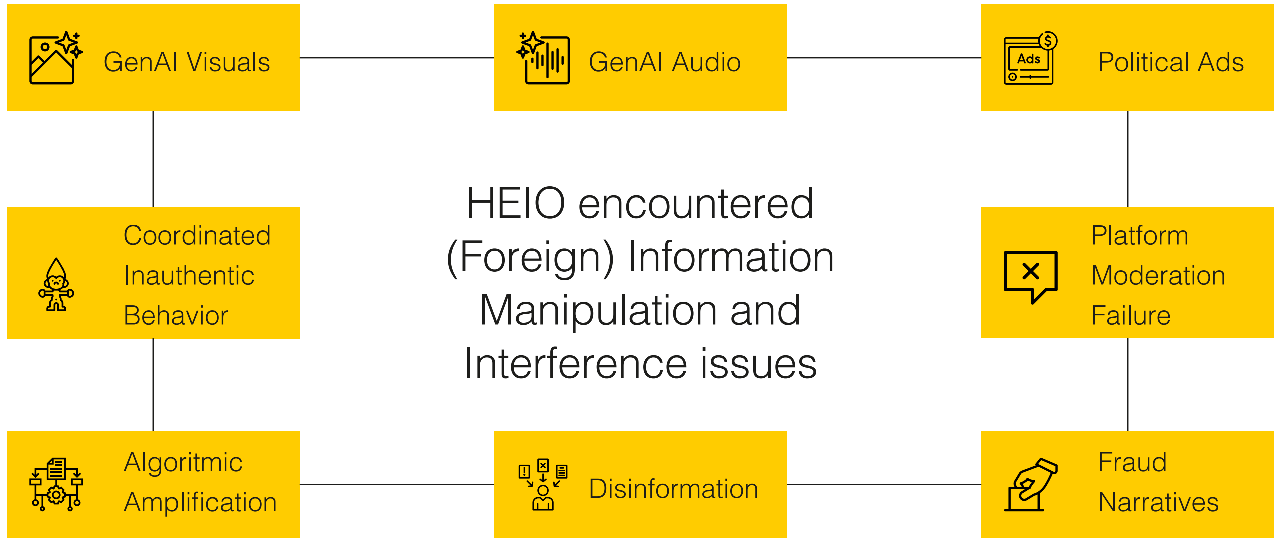

GenAI Visuals

GenAI Audio

Political Ads

Coordinated Inauthentic Behavior

HEIO encountered (Foreign) Information Manipulation and Interference issues

Platform Moderation Failure

Algoritmic Amplification

Disinformation

Fraud Narratives

GenAI Visuals
Imagary (illustrations, photos, memes and videos) created using Generative Artificial Intelligence

GenAI Audio
Music and songs created using Generative Artificial Intelligence

Political Ads
Payed advertisements on platforms by political actors or with political content

Platform Moderation Failure
Failure of enforcement of platform content moderation policies

Coordinated Inauthentic Behavior
Troll armies post content or manipulating engagement

Algoritmic Amplification
Bias or manipulation of search or recommendor algorithms in social media feeds

Disinformation
Intentional spread of misleading content with the aim to disrupt democratic processes

Fraud Narratives
Casting doubt over the election integrity, without proof

GENERATIVE ARTIFICIAL INTELLIGENCE

The 2024 "super election year", with over 60 elections across the globe, was widely anticipated as the first major test of generative AI's impact on democracy. Media outlets and researchers warned of an impending AI disinformation apocalypse. Yet early investigations found a striking disconnect between theoretical potentiality and actual observed impact. Wired, for instance, documented only 78 confirmed examples of deepfakes worldwide across all these elections. Against this backdrop, many concluded that GenAI's electoral influence had been overstated.

GENAI CONTENT

1.3 %

852

GENAI POSTS IDENTIFIED

Rather than a single game-changing technology that decides elections, GenAI tools seem to mostly intensify existing dynamics: drastically lowering the cost of content creation, enabling communication strategies at scale, and adding additional means for manipulation to an already fragile information environments. With 852 posts including GenAI, our CampAIgn Tracker dashboard identified substantially more instances of AI-generated political content than previous monitoring efforts had detected elsewhere.

This is not because the Netherlands is unique, but because our tool tracked content systematically and broadly, capturing not just deepfakes but any synthetic visuals used in political communication during the Dutch election campaign. While the overall share of AI content remains relatively modest at 1.3% of monitored posts, this figure obscures more concerning trends: the reach, engagement patterns, and strategic deployment of that content by specific actors.

We did not observe an AI apocalypse, but something perhaps more insidious: the normalization of synthetic content as a routine campaign tool, deployed most aggressively by actors operating at the margins of acceptable political discourse. GenAI is not replacing traditional campaigning, but it is amplifying the capacities of those willing to push boundaries, enabling them to do far more with far less.

Two distinct patterns of visual AI adoption emerged in our data. Smaller parties with limited resources showed higher rates of AI usage, likely leveraging the technology to compensate for their resource constraints. Meanwhile, parties with fewer normative constraints, frequently positioned on the far right, as well as satirical and parody accounts, showed elevated adoption rates, using AI-generated content to provoke and shock audiences, testing the boundaries of acceptable political discourse. Well-resourced, established parties remained more existing routines since it is likely they would face greater scrutiny and reputational risks from perceived misuse.

In terms of how GenAI was deployed, two main usage patterns stand out. Roughly half of the images function as generic "stock" visuals: people in the background or foreground, volunteers, landscapes and other decorative motifs (see Figure 1). The remaining half consists of explicitly persuasive material designed to shape political attitudes, which, in the case of far-right actors, frequently relies on racist imagery.

Regardless of purpose, most of the GenAI content we tracked circulated without any disclosure: only 32% of AI-generated posts carried a label indicating synthetic origin. When labeling did occur, it was predominantly displayed via platform labels rather than disclosure by the posting accounts in text or visual of the post.

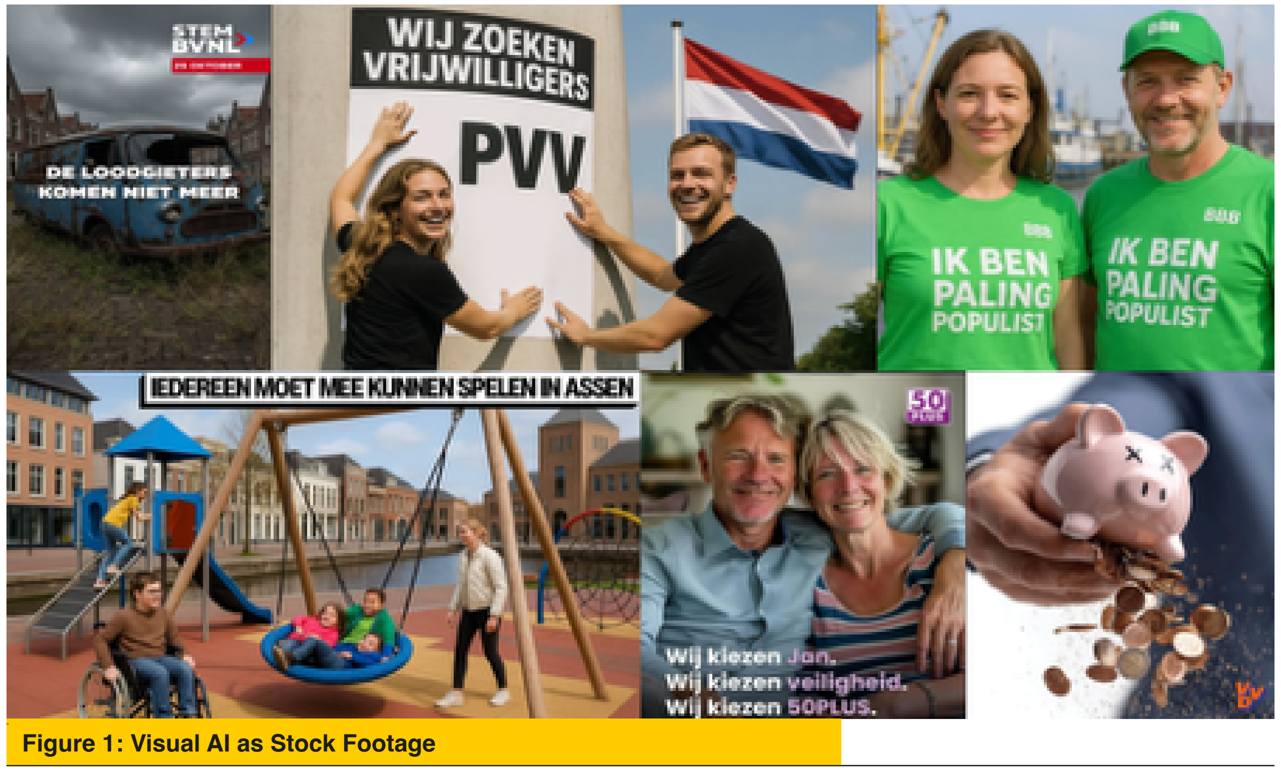

Figure 1: Visual AI as Stock Footage

PVV GENAI FACEBOOK PAGE

The most successful GenAI powered campaign we observed appeared on the Facebook page "Wij doen GEEN aangifte tegen Geert Wilders" ("We will not press charges against Geert Wilders"), which became the most popular political Facebook page in the Netherlands with sometimes over one million views per day, 74 million views between June and October 2025 in total. In December 2024, reporting by the Groene Amsterdammer already revealed that it was two PVV Members of Parliament who actively created and shared the AI-generated content on this page, which functioned as what was described as an AI-powered hate machine.

In our data, this page emerged as the single most engaged political Facebook page in the Netherlands, consistently outperforming official party accounts across the spectrum and the total engagement that Geert Wilders was receiving. The page produced a constant stream of AI-generated images targeting minorities, migrants, and political opponents, including Frans Timmermans and Henri Bontebal, often blending cartoonish aesthetics with dehumanising visuals.

The use of GenAI made it trivial to produce sometimes dozens of new images per day, lowering the cost of experimentation and allowing administrators to quickly double down on formats and themes that generated

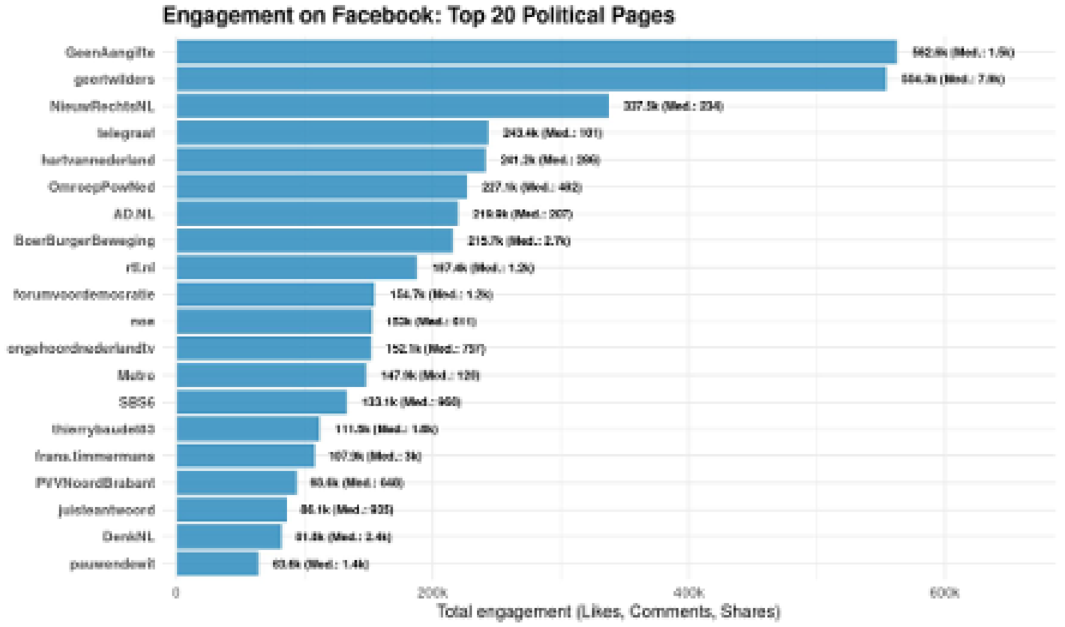

Data: CampAIgn Tracker (campaigntracker.nl). In total we track 1.2k political pages on Facebook between October 1st to October 19th 2025.

Figure 2: Top 20 Most-Engaging Political Facebook Pages

CampAIn Tracker outcomes were subsequently published in Groene Amsterdammer and Volkskrant

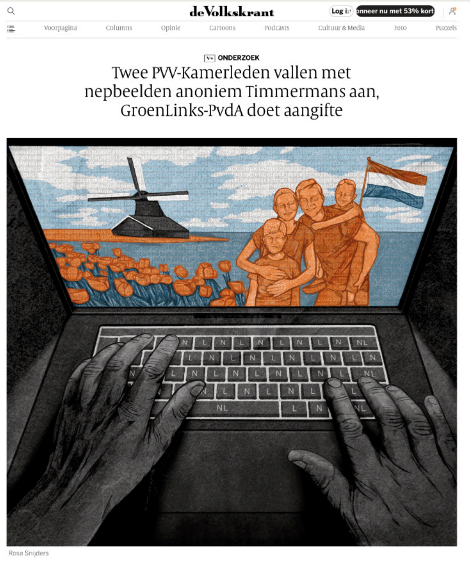

deVolkskrant

Log in | onneer nu met 53% kort

Voorpagina Columns Opinie Cartoons Podcasts Cultuur & Media Foto Puzzels

V+ ONDERZOEK

Twee PVV-Kamerleden vallen met nepbeelden anoniem Timmermans aan, GroenLinks-PvdA doet aangifte

Rosa Snijders

When a follow-up investigation by the Volkskrant found dozens of death threats against Frans Timmermans on the page, GroenLinks-PvdA filed a criminal complaint against the PVV members over the AI-generated images and the death threats circulating under them. The page was subsequently deleted.

This demonstrated that exposure and accountability can have impact, but only because the story was picked up again during the campaign period when media attention was high. The underlying platform moderation failure that allowed the page to operate for months, accumulating massive reach while systematically violating Meta's stated policies on hate speech and violent content, was never addressed. Meta took no action based on user reports or policy enforcement; the platform responded only after journalists did the investigative work that content moderation should have caught.

THEMATIC GENAI PATTERNS

Even beyond the case of two PVV members, our coding of confirmed AI-generated content revealed clear thematic patterns. The topics of migration was highly present, with synthetic imagery often depicting immigrants as threats, e.g. as criminals, invaders, cultural destroyers. Economic anxiety, energy (transition), and housing appeared frequently as well, often also framed through scarcity narratives positioning native Dutch citizens against immigrants competing for resources.

Visual styles clustered into recognizable genres. Nostalgic imagery evoked an idealized Dutch past, windmills, traditional dress, and homogeneous communities, and set it against the multicultural present. Often, it also imagined a utopian future that would once again be “free of immigrants,” mirroring that idealized past. In contrast, dystopian content showed urban decay, crime, and social breakdown attributed to immigration.

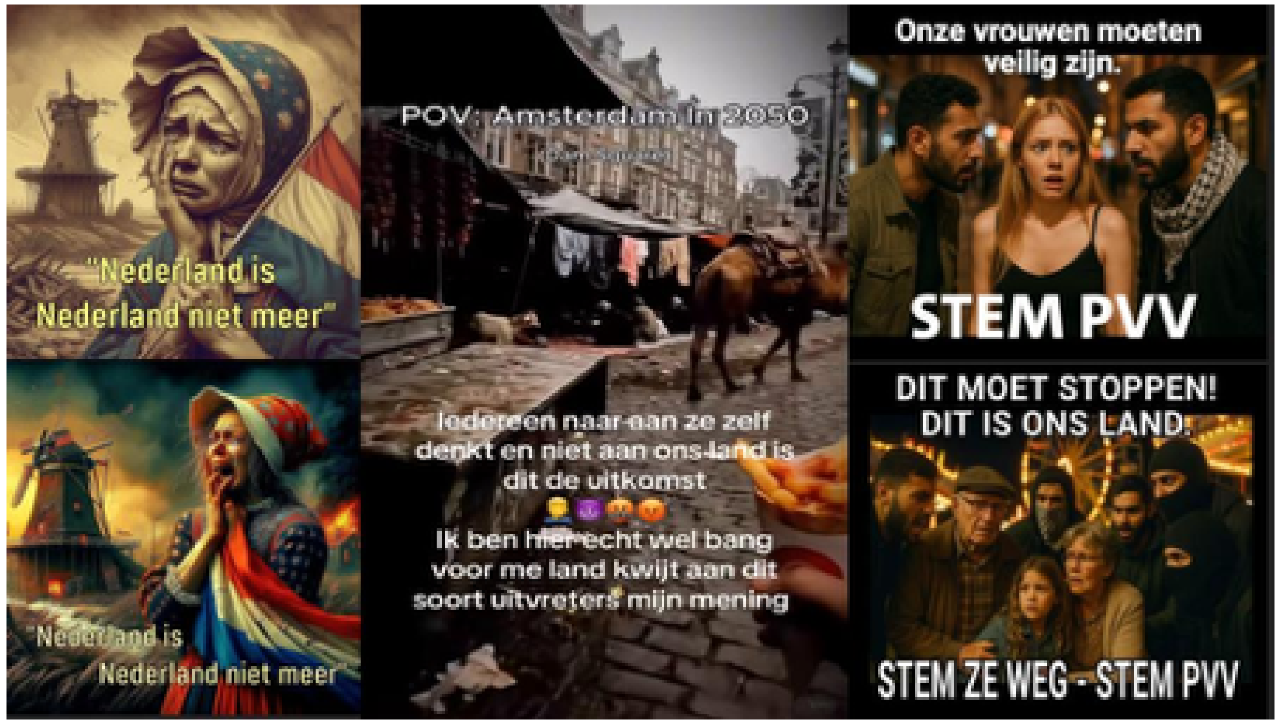
"Nederland is
Nederland niet meer"
POV: Amsterdam in 2050
Iedereen naar aan ze zelf denkt en niet aan ons land is dit de uitkomst
Ik ben hierecht wel bang voor me land kwijt aan dit soort uitvreters mijn mening
Onze vrouwen moeten veilig zijn.
STEM PVV
DIT MOET STOPPEN! DIT IS ONS LAND!
Nederland is
Nederland niet meer
STEM ZE WEG - STEM PVV

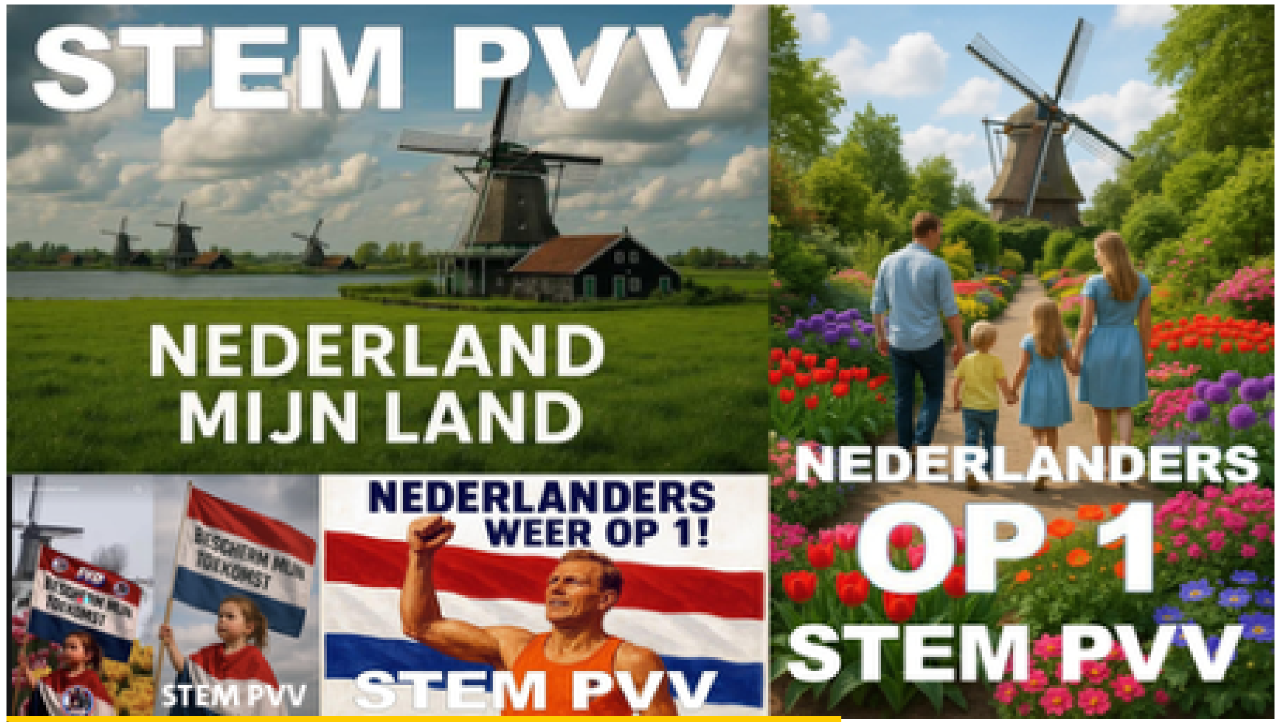

Figure 3: Nostalgic, Utopian and Dystopian AI imagery in the Dutch Election Campaign

GENAI CONTENT POLICY VIOLATIONS

Several AI-imagery we documented likely violated Dutch law. Content depicting violence against politicians showing them being physically attacked or subjected to degrading treatment potentially constitutes illegal threats under Dutch criminal code. Imagery targeting ethnic and religious minorities with dehumanizing portrayals and getting arrested may violate hate speech provisions. Beyond criminal law, much of this content involved unauthorized use of public figures' likenesses in fabricated scenarios, raising portrait rights and defamation concerns.

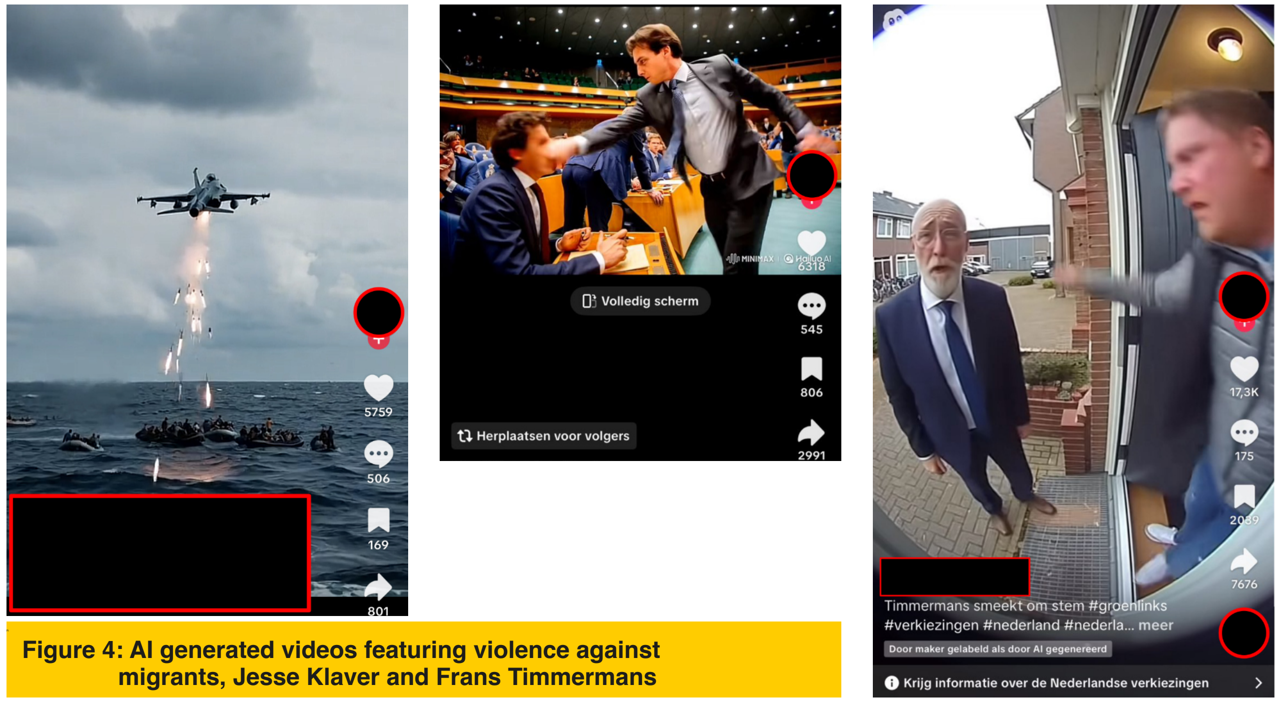
[redacted]
[redacted]

[redacted]
[redacted]

Timmermans smeekt om stem #groenlinks #verkiezingen #nederland #nederla... meer

Door maker gelabeld als door AI gegenereerd

Krijg informatie over de Nederlandse verkiezingen

Figure 4: AI generated videos featuring violence against migrants, Jesse Klaver and Frans Timmermans

The intersection of AI-generated content with humor and dark meme culture creates particular moderation challenges. Much of the most viral content occupied an ambiguous space, formatted as satire or absurdist humor, but carrying messages of genuine hostility. Defenders could claim ironic intent and this ambiguity is not accidental. It is a deliberate affordance that allows extreme content to circulate under plausible deniability. Platform content policies are poorly equipped to navigate this terrain, defaulting to inaction.

We are particularly concerned about the role of tools like OpenAI's Sora in enabling this content ecosystem. Reportedly by De Groene Amsterdammer, the Geen Aangifte page mentioned earlier mostly created its racist and hateful content using Sora despite the fact that such usage is against their terms of service. Furthermore, video generation capabilities that were experimental months ago are now producing content sophisticated enough to easily pass casual inspection. The ease with which bad actors can produce high volumes of realistic synthetic media represents a structural change in the information environment, one that current detection methods and platform policies are not equipped to address. We observed Sora-generated content depicting scenarios that would have required significant production resources to fake just two years ago, now created and distributed within minutes.

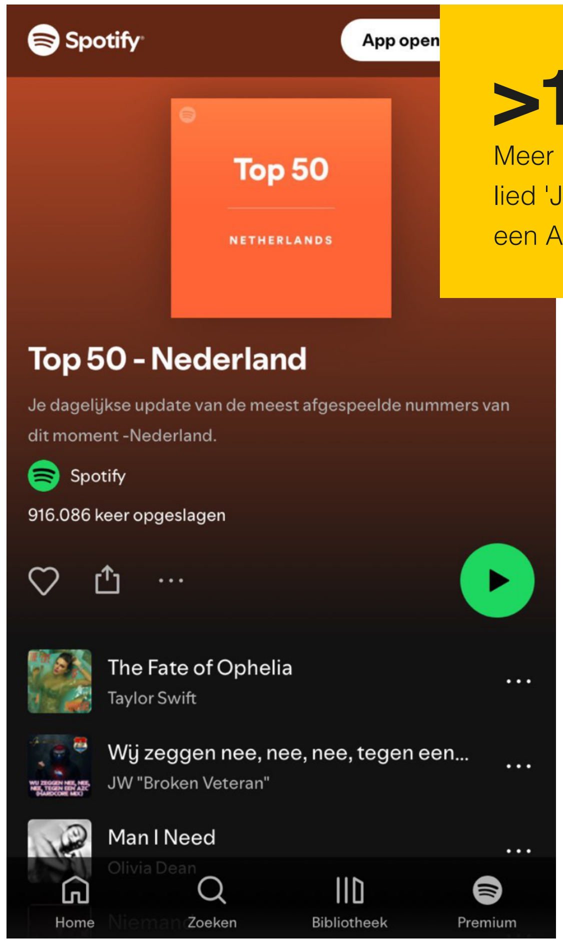

Spotify

App open

Top 50

NETHERLANDS

Top 50 - Nederland

Je dagelijkse update van de meest afgespeelde nummers van dit moment -Nederland.

Spotify

916.086 keer opgeslagen

The Fate of Ophelia
Taylor Swift

Wij zeggen nee, nee, nee, tegen een...
JW "Broken Veteran"

Man I Need
Olivia Dean

>1 miljoen streams

Meer dan een miljoen streams in 13 dagen brachten het GenAI lied 'JW "Broken Veteran" - Wij zeggen nee, nee, nee, tegen een AZC' op positie 2 in de Spotify Top50

GENAI MUSIC

An AI-generated protest song "Wij zeggen nee, nee, nee tegen een AZC" demonstrated how AI tools have lowered barriers to viral content creation. The song reached #2 on Spotify's Netherlands Top 50 and spawned over 2,500 TikTok videos by the end of October.

In November TikTok limited the use by making the song no longer “available in your country or region”, although the sound still played under videos that used the original track and plenty of remixes and reposts under a different name remained active.

JfP conducted an OSINT investigation to determine the identity of “JW Broken Veteran,” the artist behind the track. The purpose was to assess whether the artist had extremist affiliations or whether the account was operated by a non Dutch user. The identity of the individual was established. The findings indicated that neither concern applied. Based on this outcome, the identity was not published.

On 12 November all songs were taken offline from Spotify and YouTube. After Spotify confirmed that it had not removed the content, follow up research was conducted to identify the cause. The analysis concluded that the distribution partner (DistroKid) for “JW Broken Veteran” had removed the associated account. The songs were reuploaded a few days later, likely through a manual upload process. This demonstrates the difficulty platforms have with permanently removing content.

We believed it functioned as an extension of meme culture similar to the "Ausländer raus" phenomenon associated with Gigi D'Agostino in Germany.

ENGAGEMENT PATTERNS & ALGORITHMIC AMPLIFICATION

Although AI-generated visuals accounted for only a small share of political content, their posts drew about 23 times more median engagement on Facebook than non-AI posts (see Figure 5). However, on TikTok we observed the opposite: non-AI content significantly outperformed AI-generated posts and on Instagram and X AI visuals showed no statistically significant advantages to non AI content. Across platforms, AI posts from smaller and political fringe parties showed the most pronounced engagement advantages, including for example BVNL, FvD, and PVV, though for the latter this pattern is driven by the previously mentioned very succesful Geen Aangifte page.

Distinguishing organic popularity from algorithmic bias remains methodologically challenging. We cannot definitively determine whether AI content succeeds because audiences genuinely prefer it, or because platform recommendation systems inadvertently reward characteristics common to synthetic content, i.e. more provocative compositions, emotional intensity optimized through rapid iteration. Both mechanisms likely contribute. The practical implication is the same: if wielded correctly, actors deploying AI-generated content can gain competitive advantages in the attention economy of these platforms.

Non-AI vs AI-Generated Content in the 2025 Dutch Parliamentary Election

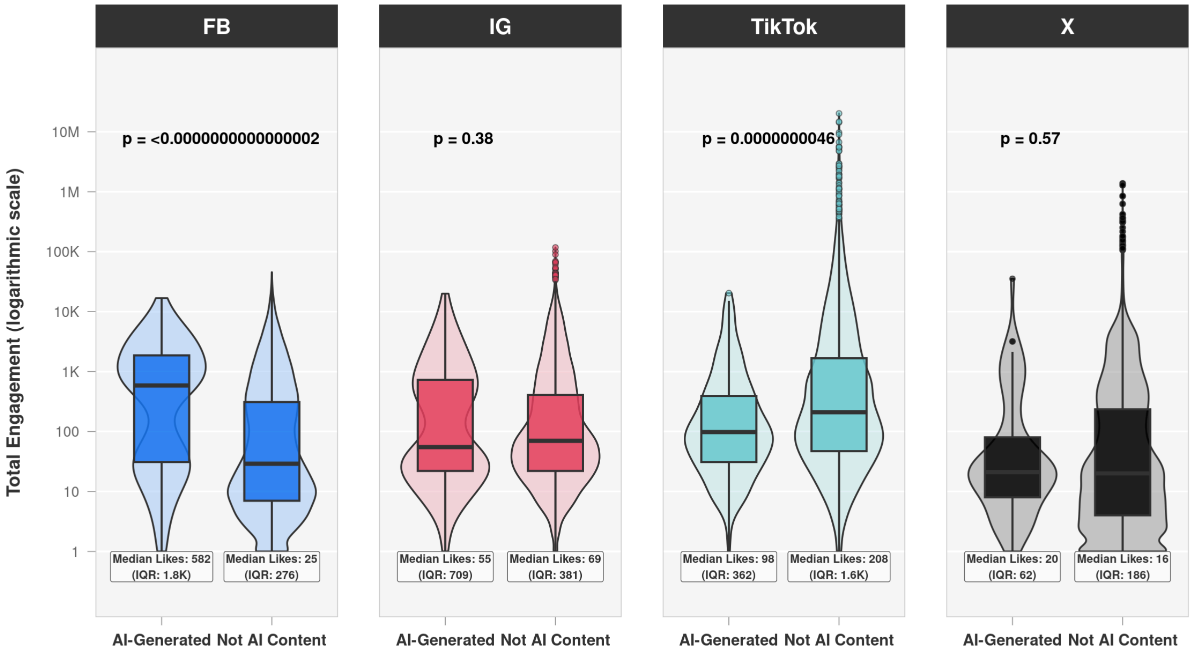

Data includes likes, comments, shares, and reactions. Statistical significance assessed via Wilcoxon rank-sum test. Source: campaigntracker.nl. Data from September 17th - October 29th 2025.

Figure 5: Comparison of Engagement between AI vs. Non-AI Visuals

COORDINATED INAUTHENTIC BEHAVIOR

Our observatory documented multiple instances of coordinated inauthentic behavior (CIB). CIB is characterized as groups of real or fake accounts work together on social media to deceive people. Often this occurs through fake engagement, such as likes and reposts, or posting similar content simultaniously. The goal is often to amplify attention to a particular issue or suppress others.

The identification of a coordinated manipulation effort often starts with finding an anomaly in the data. RTL News journalists started their search by looking into accounts who were retweeting politicians. On October 10th a number of strange accounts were detected which were retweeting mainly far right and/or polarizing content. The accounts also attracted their attention, because they had unusual names such as Anock van Dinik and Jovelyn Bryson.

Consequently, Trollrensics built specific software to make lists of people who retweeted a certain post on X. Soon hundreds of such accounts were found. In some cases the accounts used names of existing Dutch people such as the musician Andre Rieux, Madelon Vos and others. Although there are thousands of such accounts, 550 were found to repost political content.

Recently X made it possible to see locations of the accounts and it turned out that all the accounts found by RTL/Trollrensics were from Nigeria, Ghana and Ivory Coast. It is possible that Russia pays these coordinated campaigns, because in the past CNN and Guardian wrote about Russian funded troll farms in both Ghana and Nigeria. We were unable to attribute such a connection to a state actor, noting that the “influence as a service” industry is notoriously opaque.

550

Accounts on X engaging with Dutch political content from Nigeria, Ghana and Ivory Coast

Our analysis did reveal that this influence operation is not limited to election periods but operates year-round and abroad, establishing influence networks and building narratives for strategic activation during critical democratic moments.

RTL made a long article for their website, it was the most read article of the day. The discovery was also discussed on the RTL News TV broadcasts several times. Several media companies wrote about the investigation, including Volkskrant. That article had a record number of readers in a day (150,000).

The investigation by Trollrensics is still ongoing and new articles will likely be published in the nearby future. The project executed by RTL and Trollrensics took 6 weeks to complete, four people from RTL were involved and three people from Trollrensics.

Trollrensics trained several University of Amsterdam (UVA) researchers to use Trollrensics software. With a group of 3 UVA researchers and two Trollrensics employees a project was executed to analyze posts by Dutch politicians for Russian narratives. The research was finalized and will feature in a documentary by VPRO/HUMAN that will be published in the near future.

CIB ON WILDERS

On July 12, 2025, Geert Wilders made a post about how Islamic schools should be banned. This was made on his personal account. Just seven minutes later, the account "Inevitable West" made a post about this in English.

Since then, the post has been reposted in various iterations by at least 40 accounts. In total, we find the same reposted message more than 60 times on X, as well as other messages on TikTok, Threads, Facebook, and Instagram. On X alone, these posts have a reach of approximately 9 million. This shows that a larger network, largely made up of bots, contributed to the artificial spread of this statement.

Such bot networks ensure the spread and virality of misinformation and disinformation across all social media platforms.

VIEWS GAINED THROUGH CIB ON SINGLE POST

9 million

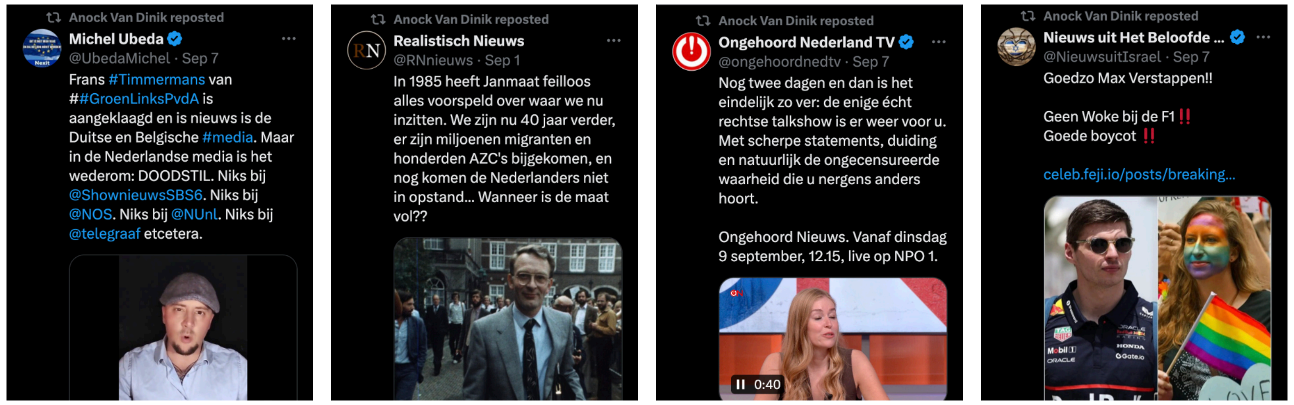

Anock Van Dinik reposted
Michel Ubeda @UbedaMichel · Sep 7
Frans #Timmermans van ##GroenLinksPvdA is aangeklaagd en is nieuws is de Duitse en Belgische #media. Maar in de Nederlandse media is het wederom: DOODSTIL. Niks bij @ShownieuwsSBS6. Niks bij @NOS. Niks bij @NUnl. Niks bij @telegraaf etcetera.

Anock Van Dinik reposted
Realistisch Nieuws @RNnieuws · Sep 1
In 1985 heeft Janmaat feilloos alles voorspeld over waar we nu inzitten. We zijn nu 40 jaar verder, er zijn miljoenen migranten en honderden AZC's bijgekomen, en nog komen de Nederlanders niet in opstand... Wanneer is de maat vol??

Anock Van Dinik reposted
Ongehoord Nederland TV @ongehoordnedtv · Sep 7
Nog twee dagen en dan is het eindelijk zo ver: de enige écht rechtse talkshow is er weer voor u. Met scherpe statements, duiding en natuurlijk de ongecensureerde waarheid die u nergens anders hoort.

Ongehoord Nieuws. Vanaf dinsdag 9 september, 12.15, live op NPO 1.

Anock Van Dinik reposted
Nieuws uit Het Beloofde ... @NieuwsuitIsrael · Sep 7
Goedzo Max Verstappen!!

Geen Woke bij de F1 ‼️
Goede boycot ‼️

celeb.feji.io/posts/breaking...

Figure 6: Reposts by Anock Van Dinik

CIB ON TIMMERMANS

At the start of our observations, we immediately stumbled upon a coordinated campaign ran from Vietnam and deployed on Frans Timmermans content on Facebook. This operation consisted of around 23,000 accounts mainly engaging through likes, thereby possibly influencing the visibility of that content.

Further investigation led us to a Vietnamese company that provides these kinds of fake social media engagements as a service.

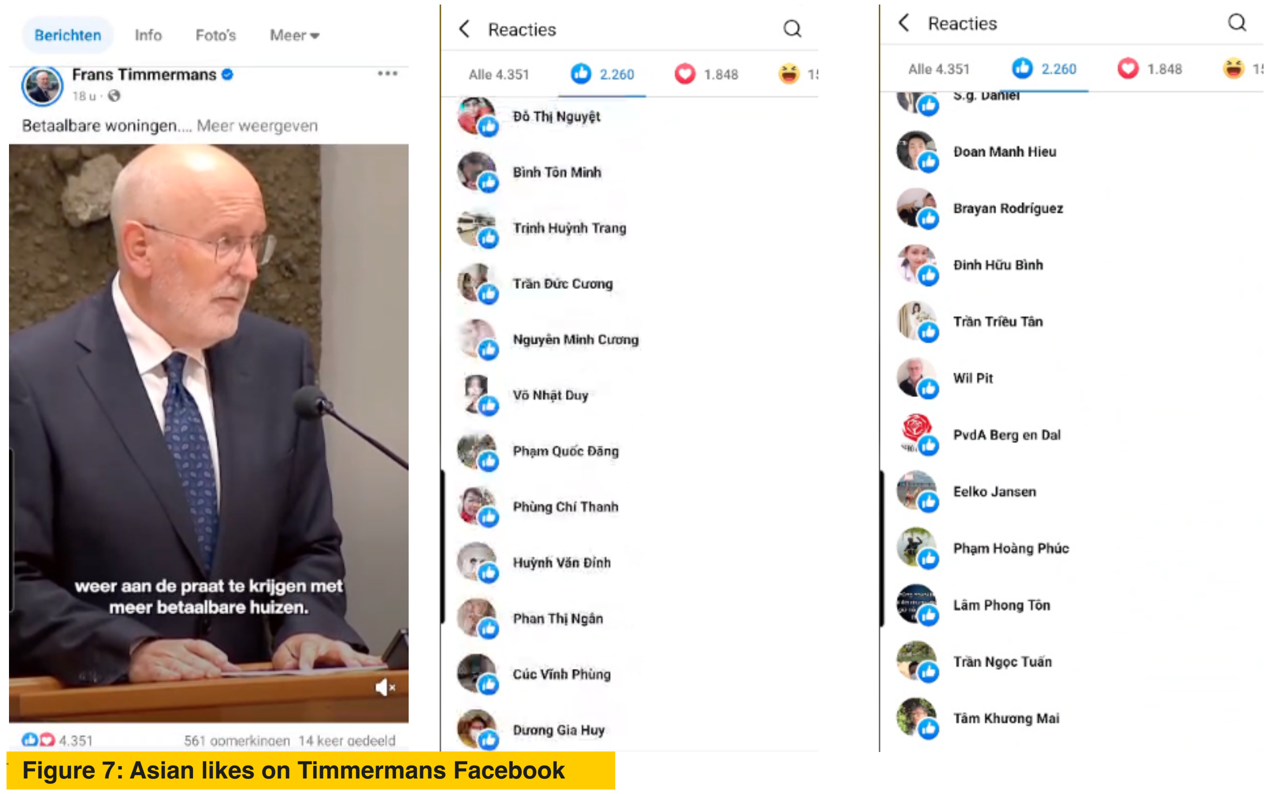

Frans Timmermans
18 u ·
Betaalbare woningen.... Meer weergeven

weer aan de praat te krijgen met meer betaalbare huizen.

4.351 561 opmerkingen 14 keer gedeeld

Đỗ Thị Nguyệt
Bình Tôn Minh
Trịnh Huỳnh Trang
Trần Đức Cương
Nguyễn Minh Cương
Võ Nhật Duy
Phạm Quốc Đăng
Phùng Chí Thanh
Huỳnh Văn Đính
Phan Thị Ngân
Cúc Vĩnh Phùng
Dương Gia Huy

S.g. Daniel
Đoan Manh Hieu
Brayan Rodríguez
Đinh Hữu Bình
Trần Triều Tân
Wil Pit
PvdA Berg en Dal
Eelko Jansen
Phạm Hoàng Phúc
Lâm Phong Tôn
Trần Ngọc Tuấn
Tâm Khương Mai

Figure 7: Asian likes on Timmermans Facebook

DISINFORMATION

We detected traditional forms of disinformation alongside the "fake" AI-generated content already mentioned in this report. Classical disinformation included false claims, f.e. about police instigating violence at demonstrations. Other misleading information with the intention to manipulate the election process, included impersonation of public figures as well as organisations, such as fake Antifa pages.

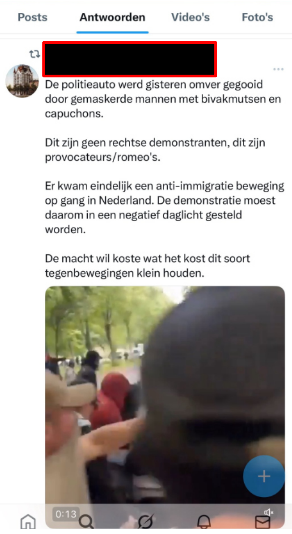

[redacted]

De politieauto werd gisteren omver gegooid door gemaskerde mannen met bivakmutsen en capuchons.

Dit zijn geen rechtse demonstranten, dit zijn provocateurs/romeo's.

Er kwam eindelijk een anti-immigratie beweging op gang in Nederland. De demonstratie moest daarom in een negatief daglicht gesteld worden.

De macht wil koste wat het kost dit soort tegenbewegingen klein houden.

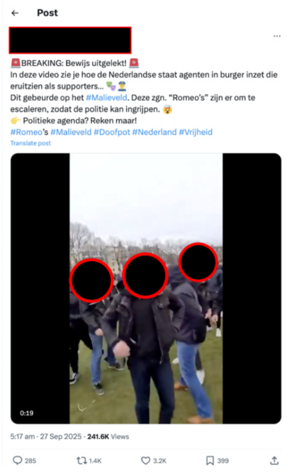

[redacted]

🚨BREAKING: Bewijs uitgelekt! 🚨
In deze video zie je hoe de Nederlandse staat agenten in burger inzet die eruitzien als supporters... 🎭👮
Dit gebeurde op het #Malieveld. Deze zgn. "Romeo's" zijn er om te escaleren, zodat de politie kan ingrijpen. 🤯
👉 Politieke agenda? Reken maar!
#Romeo's #Malieveld #Doofpot #Nederland #Vrijheid
Translate post

5:17 am · 27 Sep 2025 · 241.6K Views

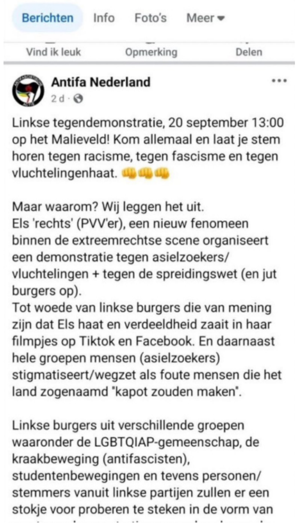

Antifa Nederland
2 d

Linkse tegendemonstratie, 20 september 13:00 op het Malieveld! Kom allemaal en laat je stem horen tegen racisme, tegen fascisme en tegen vluchtelingenhaat. 👊👊👊

Maar waarom? Wij leggen het uit.
Els 'rechts' (PVV'er), een nieuw fenomeen binnen de extreemrechtse scene organiseert een demonstratie tegen asielzoekers/vluchtelingen + tegen de spreidingswet (en jut burgers op).
Tot woede van linkse burgers die van mening zijn dat Els haat en verdeeldheid zaait in haar filmpjes op Tiktok en Facebook. En daarnaast hele groepen mensen (asielzoekers) stigmatiseert/wegzet als foute mensen die het land zogenaamd "kapot zouden maken".

Linkse burgers uit verschillende groepen waaronder de LGBTQIAP-gemeenschap, de kraakbeweging (antifascisten), studentenbewegingen en tevens personen/stemmers vanuit linkse partijen zullen er een stokje voor proberen te steken in de vorm van

Figure 8: Disinformation on X and Facebook related to the Sept 20th anti-migration demonstration. Including impersonation of Antifa, suggesting police started violence or exposing undercover agendats, using footage from a protest in Amsterdam several years ago

RUSSIAN DISINFO OPERATIONS

Known Russian information operations (f.e. Doppelgänger, Pravda) showed minimal detectable impact. We found no significant Doppelgänger patterns on Meltwater, and no traditional copycat websites. Pravda linked to official FvD channels and picked up election fraud narratives including, primarily via Telegram with minimal presence in other media.

We did notice a continuation of the foreign interference within European elections. The main actors were Russian state propaganda outlets such as RIA Novosti and RT (Russia Today or Rossiya Segodnya). Despite the blanket ban imposed on most Russian state-owned news and broadcasting companies, these outlets are still frequently used as sources for other news outlets.

False claims of manipulation of the Dutch democracy and votes was still spread online. This applied to other politically controversial subjects that could influence large numbers of people through social-media or online virality.

For example RT falsely linked the Malieveld riots to the Lisa murder case, capitalizing on untrue chaos and societal breakdown. RIA Novosti linked the results of the Dutch elections to the "deep state", claiming that the elections were falsified.

While this information was not as readily available to the public, it nevertheless strengthened the already existing election fraud narrative and suspicions.

IMPERSONATIONS & DEEPFAKES

Multiple impersonation attempts were documented, including fake Wilders accounts and a fake Wilders LinkedIn profile. A deepfake video of the King Willem Alexander circulated, alongside cheaper manipulated content of political figures including Timmermans and Jetten. While some content was clearly satirical, the Irish elections demonstrated potential real-world impact of even obvious fakes on electoral outcomes.

On TikTok we observed many impersonation accounts, featuring political figures. Some were removed, but many remain online weeks later, even though the accounts can easily be identified based on the username and/or profile picture. This suggests enforcing the platform's own policy against impersonations is not conducted on a regular basis.

https://www.tiktok.com/@geert_wilders5	https://www.tiktok.com/@geertwilders244
https://www.tiktok.com/@geert6673	https://www.tiktok.com/@geertwilders2022
https://www.tiktok.com/@wilderspvv	https://www.tiktok.com/@geert.wilders22
https://www.tiktok.com/@geertwilders211	https://www.tiktok.com/@geertwilders0870
https://www.tiktok.com/@geert.wilders56	https://www.tiktok.com/@geertwilders137
https://www.tiktok.com/@geertwilders021	https://www.tiktok.com/@geertje123452
https://www.tiktok.com/@geertwilders250	https://www.tiktok.com/@geertwilders5963
https://www.tiktok.com/@geertwildersclips	https://www.tiktok.com/@geertwilders91
https://www.tiktok.com/@geertwilders1939	https://www.tiktok.com/@geertwilders.34
https://www.tiktok.com/@geertwilders12344	https://www.tiktok.com/@geertwilders8
https://www.tiktok.com/@geertwildersfan3	https://www.tiktok.com/@geertwilders063
https://www.tiktok.com/@geertwilders.nl	https://www.tiktok.com/@geertwilders__
https://www.tiktok.com/@geertwilders65	https://www.tiktok.com/@geertwilders82
https://www.tiktok.com/@geertwilderspvvedits	https://www.tiktok.com/@geertwilders927
https://www.tiktok.com/@geert.wilders714	https://www.tiktok.com/@geertwilminder
https://www.tiktok.com/@g.wilders0332	https://www.tiktok.com/@geertwilders
https://www.tiktok.com/@geert.wilders97	https://www.tiktok.com/@geertwilders32
https://www.tiktok.com/@geert.wilders73	https://www.tiktok.com/@geertwilders0
https://www.tiktok.com/@geertjewilders123	https://www.tiktok.com/@geertwilders67
https://www.tiktok.com/@geertwilders.g	https://www.tiktok.com/@geertwilders050
https://www.tiktok.com/@geertwilders87	https://www.tiktok.com/@geertwilders97
https://www.tiktok.com/@geertwilders176	https://www.tiktok.com/@geertwilders41
https://www.tiktok.com/@geertwilders9	https://www.tiktok.com/@ikbengeertwilders
https://www.tiktok.com/@geertwilders340	https://www.tiktok.com/@geertwilders52
https://www.tiktok.com/@geertwilders97	https://www.tiktok.com/@geertwilders15
https://www.tiktok.com/@geertwilders41	https://www.tiktok.com/@_.geertwilders._
https://www.tiktok.com/@ikbengeertwilders	https://www.tiktok.com/@geertwilders48
https://www.tiktok.com/@geertwilders52	https://www.tiktok.com/@geertwilders17

Table 1: Geert Wilders impersonation accounts on October 28th, 2025 on TikTok

DEMONSTRATION RELATED DISINFORMATION

We observed that offline events such as protests and other incidents frequently served as catalysts for disinformation. In the Antifa example, the disinformation contributed to the offline event itself, while the next two examples found by JfP show how disinformation can serve to change public perception of past events.

Following the threat against Frans Timmermans on October 12, a conspiracy narrative quickly emerged alleging the perpetrator was an undercover police agent (a so-called 'Romeo'). This theory relied on a blurry picture to claim the suspect was carrying a service weapon. Notably, the image used was unsharp, possibly on purpose, as a clear image would have revealed the object was actually a waist bag.

The spread was significant: the initial post on X generated 106,600 views, followed by a Facebook post with 395 shares within a day. A subsequent post on X "confirming" the weapon theory gained another 123,600 views.

Disinformation also targeted the "Rode Lijn" demonstration in Amsterdam on October 5 regarding the situation in Gaza. False claims circulated that a participant was carrying a sign with the text "May everyday be October 7."

Although this photo was not fake, the sign was from a very different protest in Bonn, Germany. This narrative achieved a reach of 277,014 within just a few hours.

The story was picked up by several media outlets who later deleted their statements, but not before the narrative had spread widely.

It was further complicated by Grok (X's AI chatbot), which initially misattributed the fabricated image to

ELECTION FRAUD NARRATIVES

Election fraud conspiracy theories emerged as anticipated by our team, amplified by some political figures including Geert Wilders. X saw hockey-stick growth curves in posts related to fraud claims, with one account systematically collecting "my vote wasn't counted" posts. Conspiracy theories referenced historical incidents like software pentests by a company whose owner is linked to the D66 party from 2020, and specific conspiracy theorists gained some prominence questioning the integrity of the election council.

Notably, institutional actors such as the Kiesraad, municipalities, and academic experts responded rapidly to counter false narratives when these emerged in professional media. This demonstrated that prepared rapid response mechanisms can effectively limit damage from fraud claims. We noticed that these narratives were picked up by Russian propaganda outlets like Pravda, with minimal measurable impact in Dutch discourse.

YouTube data collection and analysis of top search results revealed small-scale spread of fraud narratives through video content. While only nine videos concerned fraud allegations and were primarily posted post-election, these videos had high levels of comments and likes engagement compared to general election-related videos surfaced by the search algorithm. While there does not appear to be algorithmic amplification of fraud narratives, pre-existing audiences on YouTube were highly engaged in the conspiracy.

CHATBOT IMPACT

PXS noticed that the impact extended even to AI chatbots. Anthropic's Claude and OpenAI's ChatGPT began to incorporate fraud narratives into their responses when asked about the integrity of Dutch elections, but quickly started to quote the institutional response as well.

Claude and ChatGPT began to incorporate fraud narratives the their responses

ELECTION FRAUD NARRATIVES ON X

Similarly, JfP also anticipated a rise in popularity for election fraud conspiracies around the time of the results being released. Therefore, during the elections and right after, JfP has monitored the use of tags "Verkiezingsfraude" (voiting fraud) and "Stemfraude" (votefraud) on Twitter/X.

We observed that these tags clearly spiked in popularity right on the day of the election and maintained virality in the following days. The popularity was likely exacerbated by the viral posts made by Geert Wilders, who claimed that the elections and the exit polls were fraudulent. Wilders' conspiracy generated more discussions, amplifying this false narrative.

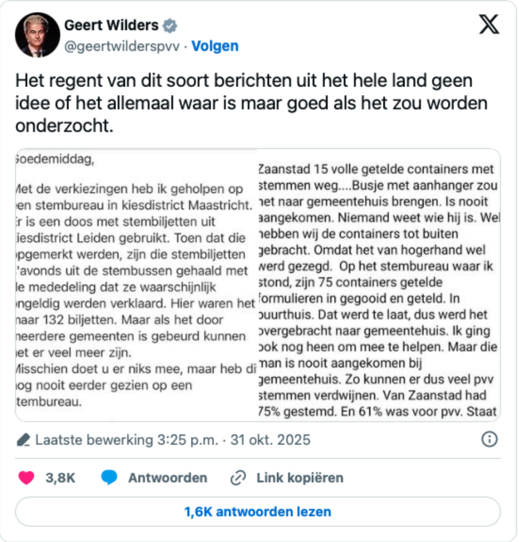

Geert Wilders
@geertwilderspvv · Volgen

Het regent van dit soort berichten uit het hele land geen idee of het allemaal waar is maar goed als het zou worden onderzocht.

ioedemiddag,

Aet de verkiezingen heb ik geholpen op en stembureau in kiesdistrict Maastricht. :r is een doos met stembiljetten uit iesdistrict Leiden gebruikt. Toen dat die pgemerkt werden, zijn die stembiljetten 'avonds uit de stembussen gehaald met le mededeling dat ze waarschijnlijk ngeldig werden verklaard. Hier waren het naar 132 biljetten. Maar als het door neerdere gemeenten is gebeurd kunnen iet er veel meer zijn.
Aisschien doet u er niks mee, maar heb di iog nooit eerder gezien op een tembureau.

Zaanstad 15 volle getelde containers met stemmen weg....Busje met aanhanger zou net naar gemeentehuis brengen. Is nooit aangekomen. Niemand weet wie hij is. Wel nebben wij de containers tot buiten gebracht. Omdat het van hogerhand wel werd gezegd. Op het stembureau waar ik stond, zijn 75 containers getelde formulieren in gegooid en geteld. In buurthuis. Dat werd te laat, dus werd het overgebracht naar gemeentehuis. Ik ging ook nog heen om mee te helpen. Maar die man is nooit aangekomen bij gemeentehuis. Zo kunnen er dus veel pvv stemmen verdwijnen. Van Zaanstad had 75% gestemd. En 61% was voor pvv. Staat

Laatste bewerking 3:25 p.m. · 31 okt. 2025

Figure 9: Geert Wilders raising questions on election integrity

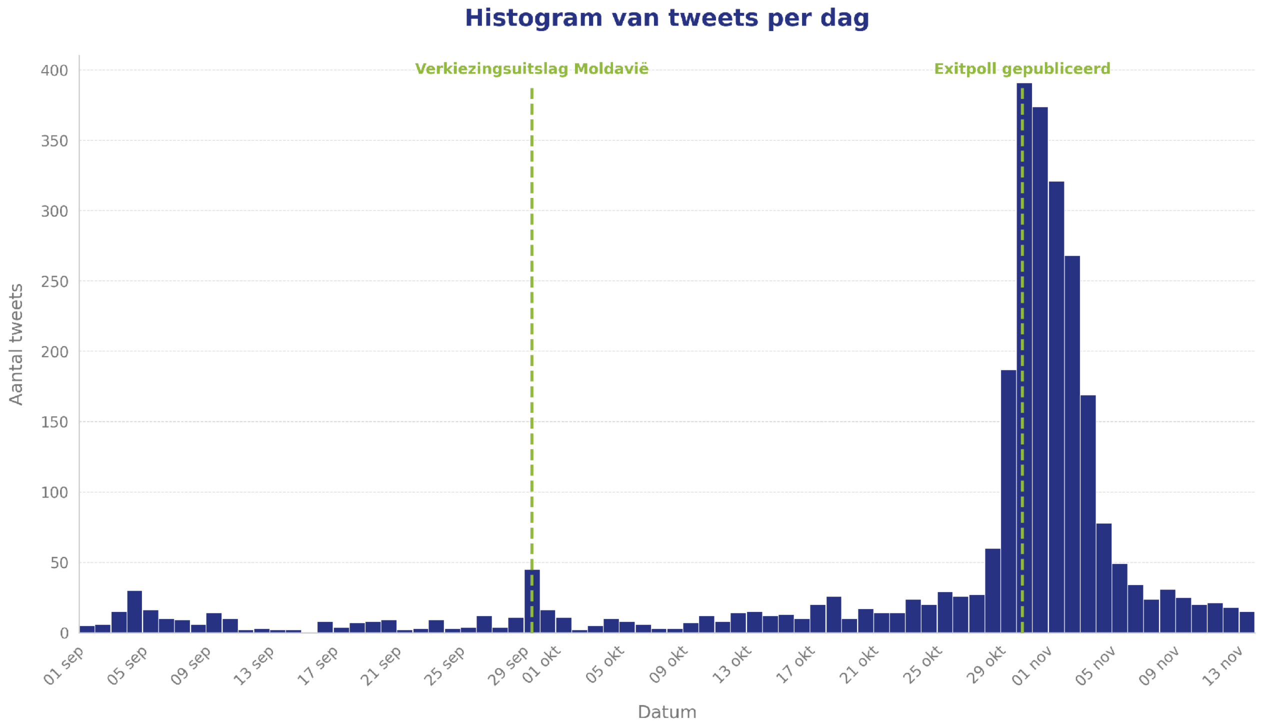

Figure 10: Number of X posts mentioning fraud related terms

COORDINATED ELECTION MONITORING

Citizens appeared to coordinate election monitoring through X. By sharing pictures of their voting ballots and stating the exact polling station, with the intention to check afterwards in the polling station report whether the votes was counted.

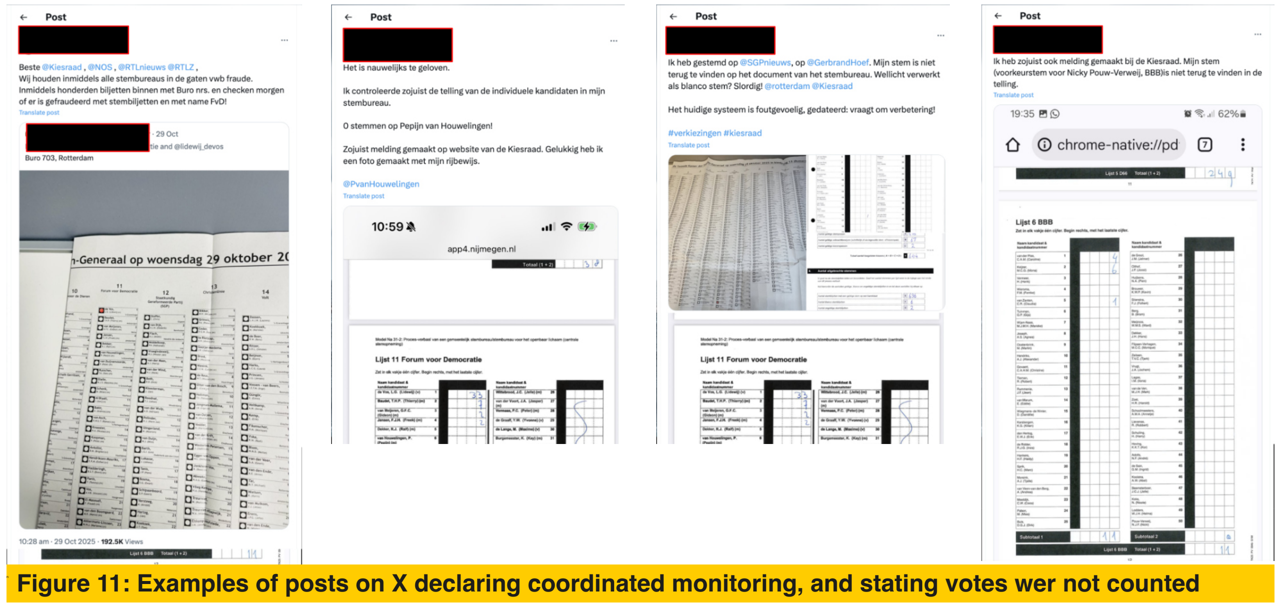

Post

[redacted]

Beste @Kiesraad , @NOS , @RTLnieuws @RTLZ ,
Wij houden inmiddels alle stembureaus in de gaten vwb fraude. Inmiddels honderden biljetten binnen met Buro nrs. en checken morgen of er is gefraudeerd met stembiljetten en met name FvD!

Translate post

[redacted] · 29 Oct
tie and @lidewij_devos
Buro 703, Rotterdam

10:28 am · 29 Oct 2025 · 192.5K Views

Post

[redacted]

Het is nauwelijks te geloven.

Ik controleerde zojuist de telling van de individuele kandidaten in mijn stembureau.

0 stemmen op Pepijn van Houwelingen!

Zojuist melding gemaakt op website van de Kiesraad. Gelukkig heb ik een foto gemaakt met mijn rijbewijs.

@PvanHouwelingen

Translate post

10:59
app4.nijmegen.nl

Lijst 11 Forum voor Democratie

Post

[redacted]

Ik heb gestemd op @SGPnieuws, op @GerbrandHoef. Mijn stem is niet terug te vinden op het document van het stembureau. Wellicht verwerkt als blanco stem? Slordig! @rotterdam @Kiesraad

Het huidige systeem is foutgevoelig, gedateerd: vraagt om verbetering!

#verkiezingen #kiesraad

Translate post

Lijst 11 Forum voor Democratie

Post

[redacted]

Ik heb zojuist ook melding gemaakt bij de Kiesraad. Mijn stem (voorkeurstem voor Nicky Pouw-Verweij, BBB)is niet terug te vinden in de telling.

Translate post

19:35 62%
chrome-native://pd

Lijst 5 D66 Totaal (1 + 2)

Lijst 6 BBB

Zet in elk vakje één cijfer. Begin rechts, met het laatste cijfer.

Figure 11: Examples of posts on X declaring coordinated monitoring, and stating votes wer not counted

ELECTION FRAUD NARRATIVES ON YOUTUBE

From 20 October 2025 to 5 November 2025, AIF found a total of nine videos that supported fraud allegations in the Dutch elections. These videos came from seven creators, two of whom made two videos each. Of those nine videos, only one was published in the week leading up the election, whereas the rest were published from 1 November 2025 on, after Geert Wilders claimed voter fraud on 31 October 2025. In comparison, election-related videos were posted more frequently before the election, spiked around the election, and were lowest after the election. The publication timing of the fraud allegation videos suggests that they are a response to Wilders's claims and the election results, rather than a coordinated and/or preempted disinformation campaign.

Across the videos, the fraud allegations are presented through a combination of speculative claims, anecdotes, or conspiratorial narratives.

Here, fraud is framed as both technically possible, asserting that the Dutch voting system is generally reliable but “not flawless” due to administrative errors like vote miscounts.

However, fraud is also framed as politically-motivated manipulation, escalating to systemic fraud by partisan actors allegedly manipulating ballots or vote checking software. The D66, as the winner of the elections, is portrayed as the main beneficiary and primarily blamed. Mainstream media and institutional authorities are depicted as enablers or the ones to legitimize the allegedly fraudulent outcomes. The fraud narrative reflects themes of broader anti-establishment, institutional distrust, as well as political polarization. Frequent discussions revolved around migration policies, EU influence, housing, and general state control, whereby governmental legitimacy and democratic representations are strongly criticized.

Fraud-related videos tended to have 2x more engagement than general-election videos that showed up in the search results. They had 2x as many likes and nearly 2x as many comments (Figure 12) compared to general election-related videos. However, general videos about the election had 1.3x as many views as fraud videos (Figure 12). While videos about fraud were viewed less, their viewers were much more engaged in the content.

When we examined video ranking in the search results, we found that fraud-related videos appeared in the top 10 search results around 29% of the time, whereas general-elected related videos appeared in the top 10 search results only 13% of the time. As we were searching with fraud-related queries, this is not so unexpected. However, the fraud-videos that appeared most often in the top 10 search results, often had the lowest engagement numbers of the nine fraud-videos. The videos' high rank is likely determined by their use of the words “Verkiezingen 2025” “fraude” and/or “Verkiezingsbedrog” in the title, providing an exact match for our search queries.

The difference between the high rank and low engagement levels of these videos, and low rank and high engagement of some of the others, suggests little emphasis from YouTube's algorithm on engagement metrics when surfacing results. Further, the engagement that the fraud videos received was less likely due to their ranking in the search results, but rather from a dedicated user base who already were following the creators who made fraud allegations. Therefore, YouTube's search ranking may play less of a role in spreading fraud allegations compared to the community cultivation that the platform affords around common interests and political alignments with respect to the October 2025 Dutch elections.

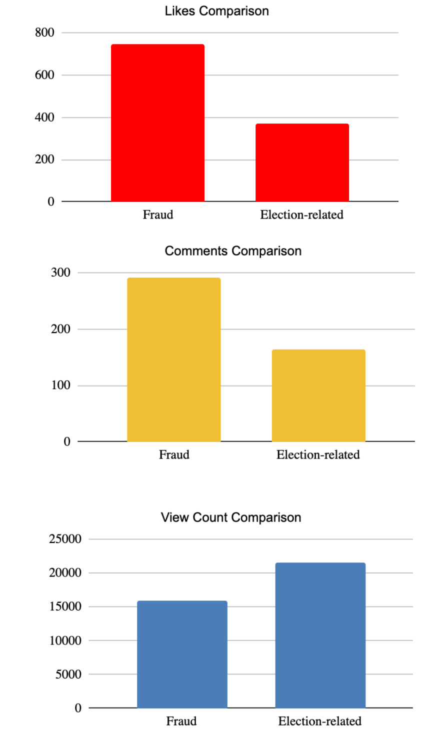

Figure 12: Likes, comments and views comparison between fraud and election-related videos

POLITICAL ADVERTISING

This election was the least transparent in recent Dutch history regarding political advertising. This had to do with the response of tech platforms to new EU regulation on transparency and targeting of political advertising (TTPA).

Meta enforced the EU regulation requiring political ad transparency early (October 6 instead of October 10), yet dozens of ads continued running past the deadline, including from official party accounts. The voluntary self-reporting system for digital political advertising in digital media other than social platforms proved completely ineffective. D66 was entirely missing from self-reported data, and databases contained no actual ad content (images, videos, texts), only vague descriptions.

TikTok and Snapchat ad databases were particularly inadequate. X proved the worst platform to monitor for political advertising. The search engine in the interface is broken. Researchers identified four broad categories of political ads: genuine political advertising, commercially-motivated political ads, personal political ads and scam ads.

We would also like to flag how political actors can potentially circumvent regulatory transparency. We observed for example how news blog ads with political messages continue to be able to run ads.

The lack of transparency means researchers and citizens have no visibility into money flows in digital political campaigns. Ads could theoretically be used to collect engagement data on specific populations for later targeting elsewhere, enabling micro-targeting without oversight.

These were the least transparent elections in terms of political advertising in recent history

FAKE POLITICAL PRODUCTS

Another example of lack of transparency and disclosure was observed when examining the Meta ad repository. Under the commercial advertisements category, we were able to find a few accounts that posted sales advertisements under this category.

However, among these ads, we noticed a trend of political slogans or other ideologically charged content pasted on T-shirts and other types of merchandise.

While it is not yet a spread phenomenon, it should definitely be flagged, as it seems that users can create false listings for non-existent products with the use of AI in order to influence the elections. However, this type of approach to false content may not aim to influence the political process, but rather it could also simply be done for financial gain.

Nevertheless, if platforms do not ensure adherence to the guidelines, it is possible that commercial ads will become a legitimate avenue for users to circumvent bans on specific advertising categories and continue spreading political messages without adequately disclosing and categorising them as such.

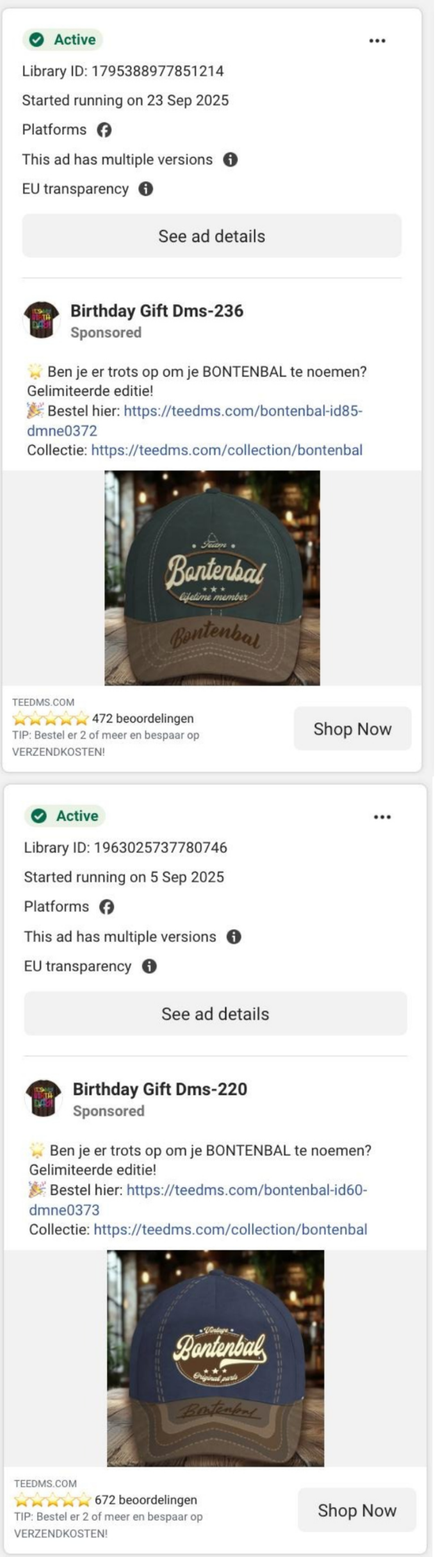

Active

Library ID: 1795388977851214

Started running on 23 Sep 2025

Platforms

This ad has multiple versions

EU transparency

See ad details

Birthday Gift Dms-236
Sponsored

🌟 Ben je er trots op om je BONTENBAL te noemen? Gelimiteerde editie!
🎉 Bestel hier: https://teedms.com/bontenbal-id85-dmne0372
Collectie: https://teedms.com/collection/bontenbal

TEEDMS.COM
⭐⭐⭐⭐⭐ 472 beoordelingen
TIP: Bestel er 2 of meer en bespaar op VERZENDKOSTEN!

Shop Now

Active

Library ID: 1963025737780746

Started running on 5 Sep 2025

Platforms

This ad has multiple versions

EU transparency

See ad details

Birthday Gift Dms-220
Sponsored

🌟 Ben je er trots op om je BONTENBAL te noemen? Gelimiteerde editie!
🎉 Bestel hier: https://teedms.com/bontenbal-id60-dmne0373
Collectie: https://teedms.com/collection/bontenbal

TEEDMS.COM
⭐⭐⭐⭐⭐ 672 beoordelingen
TIP: Bestel er 2 of meer en bespaar op VERZENDKOSTEN!

Shop Now

Figure 13: Screenshot from the Meta Ad library featuring the name Bontenbal

META'S RESPONSE TO POLITICAL ADS

In their investigation, AIF found that only nine pages were official political accounts posting political advertising in breach of the new Meta's policy.

Upon reviewing the nine ads posted by those accounts, only one was still live (Figure 14) which is from a local chapter of the GL PvDA party. Meta had removed the remaining eight advertisements likely between the time of our data collection and content review.

When it comes to finding and removing political ads posted by a political party or candidate, Meta appears seemingly successful. This does not, however, account for ads posted by entities not affiliated with a candidate or political party.

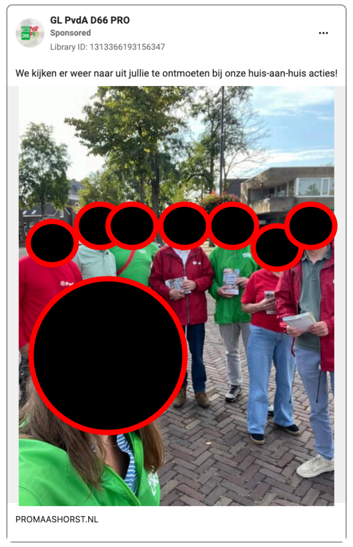

GL PvdA D66 PRO
Sponsored
Library ID: 1313366193156347

We kijken er weer naar uit jullie te ontmoeten bij onze huis-aan-huis acties!

PROMAASHORST.NL

Figure 14: The only remaining official political campaign ad on Meta

INFLUENCER MARKETING

Influencer marketing platforms like LinkPizza included categories for law, governance, and politics, suggesting infrastructure exists for paid political influence campaigns, as it was shown during the Romanian elections.

However, these platforms provided no transparency reports about political activities during the election. Influencer endorsements of political content occurred but lacked systematic disclosure, making it impossible to distinguish paid promotion from organic support. This was further made more difficult by the number of influencers posting potentially undisclosed political content through LinkPizza or similar platforms or collaborations assignments.

Due to the lack of transparency, conducting a proper investigation would need these influencers or social media personalities to be individually followed and their posts manually monitored. This was a task which fell outside of the scope of this research due to capacity and resource limitations.

SPENDING & TARGETING

Before Meta and Google's political ad bans took effect on October 10, 2025, Dutch political parties spent approximately €399k on Meta platforms, running 1,3k advertisements. Forum voor Democratie was the largest spender at €121, followed by D66 at €63k and GL-PvdA at €39k.The targeting data (see: favstats.github.io/nl25) reveals distinct strategic approaches across parties.

Custom Audiences (using parties' own voter/platform data) was the most prevalent method, consuming approximately 14% of total targeting spend, followed by postal code targeting at 10% and lookalike audiences at 9%. Interest-based targeting showed clear ideological patterns: GL-PvdA focused on users interested in social movements, journalism, and the arts, DENK targeted audiences interested in television entertainment, Arabic pop music, and African cuisine, while D66 targeted music streaming users.

These practices raise concerns about potential circumvention of TTPA transparency requirements, as the regulation prohibits targeting based on sensitive personal data like ethnicity, yet interest categories such as "couscous" or "Arabic pop music" may function as proxies for such characteristics.

Following the political ad ban, spending shifted to traditional media channels tracked through politiekereclame.nl, Ster, and DPG Media, totaling over €10 million (see: favstats.github.io/reclamer).

However, the self-reporting system proved inadequate. The three largest campaigns without channel information (D66 at €2.4 million, BBB at €400,000, and JA21 at €400,000) represent over €3.2 million in spending, and transparency databases contained only vague descriptions rather than actual ad content. The data submitted lacked the granular targeting information the TTPA was specifically designed to require: what texts, images, or videos were used, which data sources informed targeting, and how sensitive categories were avoided.

POLITICAL AD SPENDING ON META

$399.000

1.300
POLITICAL ADS ON META

D66, JA21 AND BBB SPEND

$3.2
MILLION

with only vague descriptions in stead of actual ad content according to the self-reporting system

ALGORITHMIC BIAS AND MANIPULATION

We observed activities that were aimed at manipulating the visibility of content through recommender algorithms, different from the use of troll armies. And we attempted assess biases in the recommender systems.

HIGH RANKING OF NEW ACCOUNTS

Very new accounts (maximum three days old) with minimal content (sometimes only three posts) appeared at the top of the For You Page of TikTok during the days before elections. If platforms allow this systematically, paid advertising becomes unnecessary. Actors can manipulate recommendation algorithms directly.

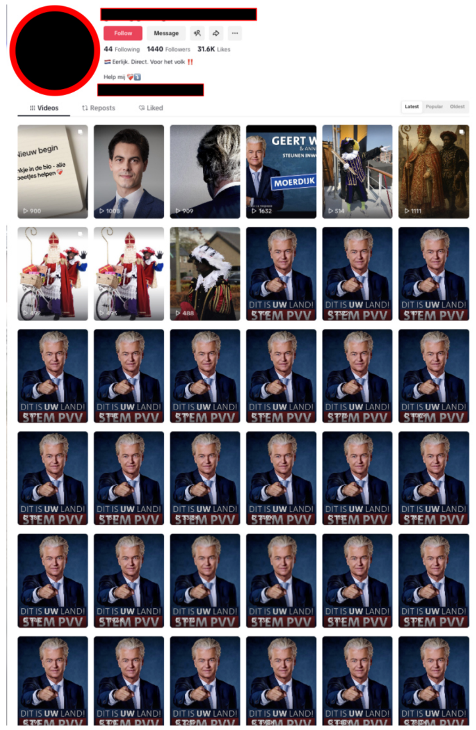
[redacted]

[redacted]

44 Following 1440 Followers 31.6K Likes

🇳🇱 Eerlijk. Direct. Voor het volk ‼️

Help mij ❤️⤵️

[redacted]

Figure 15: Account created on Oct 20th consistently posting PVV promotional videos until Oct 29th 2025, followed by divisive content on Black Pete, and other political topics

HIGH ENGAGEMENT ON FAR-RIGHT CONTENT

Far-right content consistently received higher engagement and views across platforms. Here distinguishing organic popularity from algorithmic bias remains challenging. Thus, we cannot designate this as a form of manipulation.

The lack of transparency on recommendation algorithms prevents verification of fairness. Platforms provide no data about how content is ranked, amplified, or suppressed, making it impossible to assess whether elections are conducted on level playing fields.

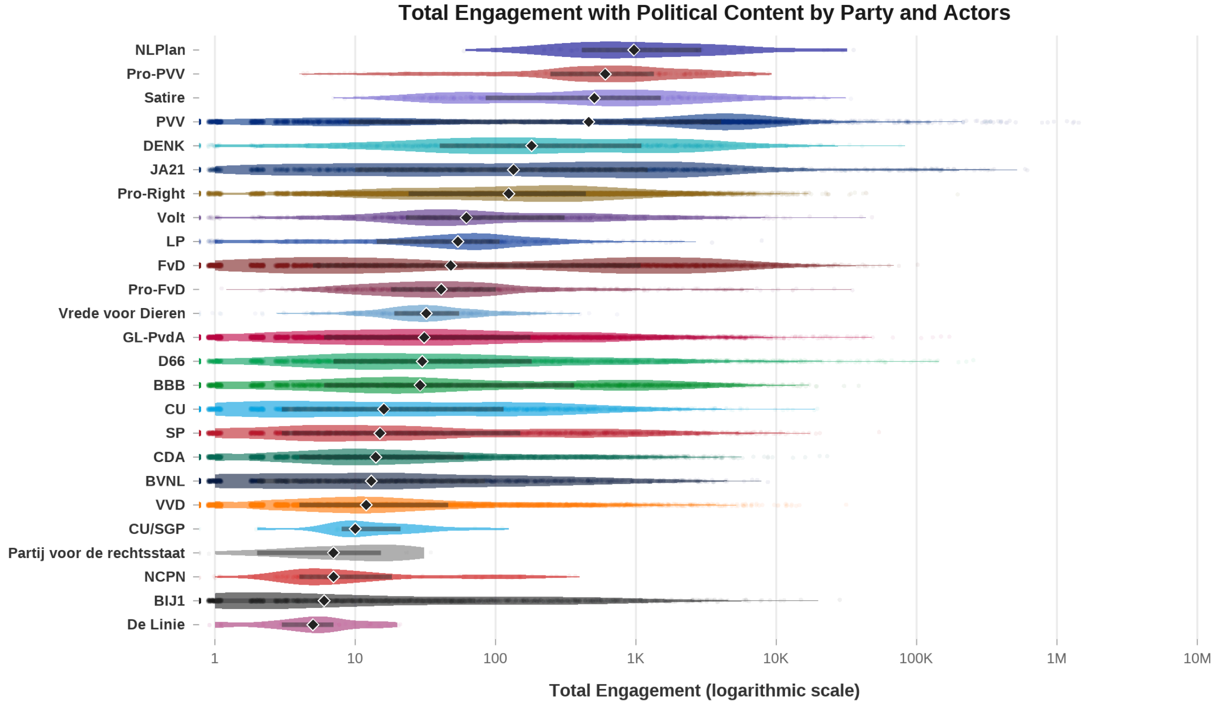

Data includes likes, comments, shares, and reactions. Source: campaigntracker.nl. Data from September 17th – October 29th 2025.

Figure 16: Total engagement with political content

PLATFORM MODERATION AND INTERACTION WITH AUTHORITIES

Platform moderation during the Dutch elections failed notably. Not a single piece of content we reported through official platform reporting mechanisms was removed, even when content clearly violated stated terms of service. This included death threats, explicit racism, antisemitism, and violent imagery shared on public political pages

The only exception occurred when HEIO brought (violent) AI-generated content to media attention. TikTok responded rapidly to media coverage, indicating that public embarrassment drives platform action instead of user safety. Even the State Secretary responded publicly, but this reaction typically ended with indignation rather than structural change. Even when the same account posted another violent AI-generated video the next day.

TikTok livestreams emerged as particularly problematic accountability-free zones. Death threats, antisemitism, and hard racism flourished without intervention. DENK party streams attracted significant international audiences, with no means for us to determine whether this included foreign interference.

Platforms made soft promises about election integrity in the period leading up to the elections, but provided no accountability mechanisms. We have no means to determine whether they lived up to the promise or not.

What makes matters worse, is that the EC Rapid Response System offered no transparency about the process or outcomes. This is due to a contractual clause that comes into effect when parties join the system. This results in us researchers essentially providing free labor to platforms, while being unable to claim results. An unsustainable model that may even divert resources from regular user reports.

We also observed how moderation falls short when live events are happening. Again, the September 20th protests serve as an example. On social media numerous posts with content featuring violence circulated.

EUROPEAN COMMISSION RAPID RESPONSE SYSTEM

The Strengthened Code of Practice on Disinformation is an EU-led framework uniting online platforms, industry actors, researchers, and civil society organizations in a shared effort to reduce the spread and impact of disinformation. Originally voluntary, it has evolved into a structured system of commitments aligned with the Digital Services Act. It now has moved toward formal recognition as a DSA Code of Conduct, a shift that elevates its status and strengthens expectations for compliance and accountability.

The Code of Conduct includes 43 commitments and 128 measures across areas such as limiting the monetization of disinformation, improving political ad transparency, and supporting users and researchers. Its work is carried out through sub-groups like monetization, fact-checking, and crisis response.

The work of the Code is further supported by mechanisms designed for rapid, coordinated action in critical situations. One key mechanism is the Rapid Response System (RRS), which enables non-platform signatories, such as civil society organizations and researchers, to alert platforms quickly about disinformation, particularly during elections or other high-risk periods.

For example, during the 2024 European Parliament elections, the RRS handled 18 notifications sent to platforms including Meta, YouTube, and TikTok. These notifications led to 12 instances of content or accounts being removed, 2 instances of content being labeled, and 1 case where both labeling and other mitigations were applied. Responses from platforms varied, including written, oral, or mixed formats. EDMO evaluated the system as effective in enabling timely action against disinformation. EDMO highlighted that the RRS alone can not address sustained narratives and cautioned against potential over-censorship if content removal is not carefully contextualized.

AI Forensics has been a signatory to the EU Code of Practice on Disinformation since 2023. In the context of the 2025 Dutch elections, AI Forensics participated as a designated contact point for the Code and the RRS, using the system to flag potentially harmful content. All interactions within the RRS are conducted under strict confidentiality, precluding the disclosure of specific details.

Our general impression is really positive, but due to contractual clauses we can provide no transparency about the decision-making process, the feedback about outcomes, or accountability for which content was addressed and why.

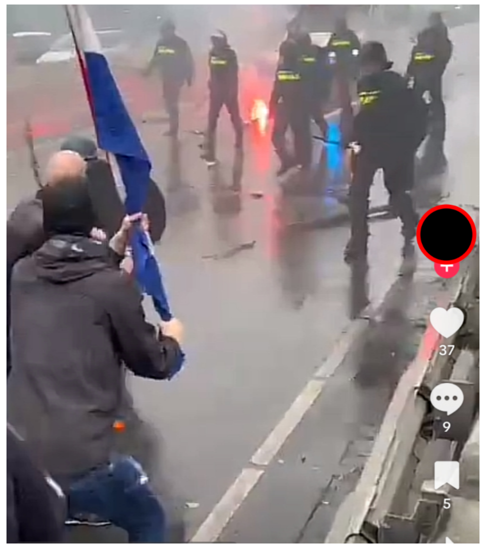

Figure 17: Screenshot from violent video on TikTok on Sept 20th

Notably, when colleagues had reported content using the standard in-app reporting function, it was not taken down through this channel. This discrepancy highlights differences in assessment methods between the RRS and standard user reporting mechanisms, raising questions regarding the prioritization and resourcing of content moderation processes. It suggests a potential need for platforms to enhance user empowerment, enabling ordinary users to participate more effectively in content moderation.

Our experience with the RRS also highlights important distinctions compared to the broader concept of Trusted Flaggers. Unlike Trusted Flaggers, which are recognized by individual platforms to report a wide range of policy-violating content on an ongoing basis, the RRS is EU-coordinated, crisis-focused, and restricted to a small circle of non-platform actors.

While this structure ensures rapid action on high-stakes issues such as elections, it could end up limiting the direct influence of ordinary users on content moderation, raising questions about equity and long-term sustainability. Additionally, participation in the RRS is voluntary and uncompensated, whereas Trusted Flaggers sometimes receive support from platforms, highlighting the resource challenges of relying on civil society actors for systemic interventions. Both systems also operate with limited transparency, making comprehensive assessment of outcomes difficult.

While the RRS represents a critical mechanism within the Code of Conduct, enabling rapid, coordinated action between civil society actors and platforms, it operates within a restricted circle of actors. Participation by non-platform signatories is voluntary and uncompensated, yet it contributes significantly to supporting the Code’s objectives.

Our experience indicates that the current model, characterized by limited transparency and reliance on voluntary engagement, may not be sustainable in the long term and risks diverting attention and resources from user-submitted reports.

Taken together, these observations suggest that while the RRS is effective in its intended scope, greater investment in routine moderation tools and user empowerment could complement its impact, ensuring more sustainable and inclusive disinformation mitigation.

OTHER AUTHORITIES

Throughout the election campaign period we interacted with several authorities.

Autoriteit Consument & Markt (ACM) – Dutch DSA Coordinator

HEIO reported findings to ACM, because the authority needs formal complaints from affected parties in order to be able to enforce regulations. This structural limitation means systematic monitoring is necessary, but it remains unclear how this is resourced now our observatory ends.

Autoriteit Persoonsgegevens (AP) – Dutch Data Protection Authority

HEIO informed AP about monitoring approach. As the AP is preparing to enforce the AI Act, our findings on the use of GenAI content and the fairness of algorithmic recommendations are relevant for the AP.

Commissariaat voor de Media (CvdM)

Informed about HEIO's monitoring methodology and findings, particularly regarding advertising transparency and media-related election integrity issues.

Ministerie van Binnenlandse Zaken (MinBZK)

State Secretary responded publicly to media coverage of HEIO findings, particularly regarding violent AI-generated content.

4

CONCLUSIONS & RECOMMENDATIONS

CONCLUSIONS

Based on our observations we draw the following conclusions.

Elections were free and fair, but under threat

The Dutch parliamentary elections of October 29, 2025, were fundamentally free and fair. We believe the legitimate outcome reflected voter preferences. However, this conclusion must be immediately contextualized. The elections took place under unprecedented digital pressure that tested the resilience of democratic institutions and exposed vulnerabilities in the information ecosystem.

Foreign actors have attempted to compromise the election integrity

Our work has exposed attempts of foreign actors to interfere with the elections. Through coordinated inauthentic behavior candidates or certain political issues received additional visibility. The strategy followed the logic of boosting already existing polarization, instead of fabricating content.

Online environment contributed to political violence

The online environment during the campaign period demonstrably contributed to real-world political violence. Death threats, violent imagery, dehumanizing content, and explicit calls for violence against politicians spread on multiple platforms without effective intervention. While direct causation is difficult to prove, the connection between online hate speech and offline violence is well-established in academic literature, and HEIO observed this dynamic during the Dutch elections.

Generative AI fundamentally changed the landscape

2024 was predicted to be the "AI super election year," but early analyses found limited AI impact globally. The Dutch elections proved this assessment premature. Generative AI has empowered bad actors to do significantly more with much less, lowering barriers to producing sophisticated propaganda, hate content, and disinformation at scale. The PVV-associated Facebook operation demonstrated that AI tools enable individual actors or small groups with the means that rival with professional political communications operations.

Most concerning: AI-generated content will become harder to detect as tools improve. We may already be at the point where GenAI content goes undetected. The current moment of obvious, detectable AI artifacts is temporary. We expect future elections will face invisible synthetic content at scale.

Transparency eroded

In our opinion this was the least transparent Dutch election in recent history with regards to political advertising. The EU regulation intended to increase political ad transparency was poorly implemented, creating less oversight than existed previously. Voluntary self-reporting proved worthless, ad libraries were incomplete and incomparable, and citizens have no ability to track money flows in digital political campaigning. This opacity benefits bad actors while hampering accountability.

Platform moderation is insufficient protection

Platform moderation during the Dutch elections was insufficiently protective. Platforms respond to public embarrassment through media coverage, not to user reports or systematic policy violations. The message is clear: terms of service exist primarily for public relations purposes, not user safety. This dynamic creates perverse incentives where researchers and activists must amplify harmful content through media channels to trigger platform response. We felt forced spreading (references to) content we actually seek to suppress, in order to trigger moderation response.

Attribution remains challenging

While we documented multiple instances of coordinated inauthentic behavior involving definitive attribution to specific actors remained elusive. This reflects fundamental challenges in attribution: actors employ plausible deniability, operations blend organic activity with manipulation.

However, lack of definitive attribution should not prevent action. Regardless of whether manipulation comes from state actors, commercial operations, or domestic extremists, the harms are real and regulatory responses must address behaviors rather than waiting for attribution.

Regulatory frameworks are inadequate

Current regulatory frameworks are structurally inadequate for protecting electoral integrity, including the DSA. Dutch authorities like ACM are powerless without formal complaints, rendering our nation dependent on a supranational body to regulate our own elections. Problematically, the mechanism of the European Commission (Rapid Response System) operates without transparency. We also experienced that DSA research access comes too late to enable real-time monitoring.

The fact that monitoring of systemic risks to democracy is being conducted by low-resourced NGOs is deeply concerning. Platforms are getting away with soft promises without accountability. This is backwards. Platforms should be required to demonstrate before elections, that their moderation systems work and are sufficient, ideally through independent certification. The burden of proof (and costs) must shift from civil society organizations to the platforms.

Threats will persist and evolve

The networks, techniques, and tactics identified during this monitoring period remain active. From the analysis of the networks we found, and from literature, we know that foreign information manipulation and interference operations are persistent, year-round activities, not limited to campaign periods. The infrastructure built for the October 2025 elections will target future general elections, municipal elections, provincial elections, and the next round of European Parliament elections.

Without structural changes to platform accountability, funding for monitoring infrastructure, algorithmic transparency, and real-time enforcement mechanisms, future elections face severe risks.

Civil society cannot bear this burden alone

The HEIO project demonstrated both the value and the limitations of civil society election monitoring. Five specialized organizations working collaboratively are able to detect patterns, expose manipulation, and influence public discourse. However, this model is not sustainable without dedicated funding, cannot scale to cover all platforms and modes of manipulation, and should not be the only way to monitor digital election

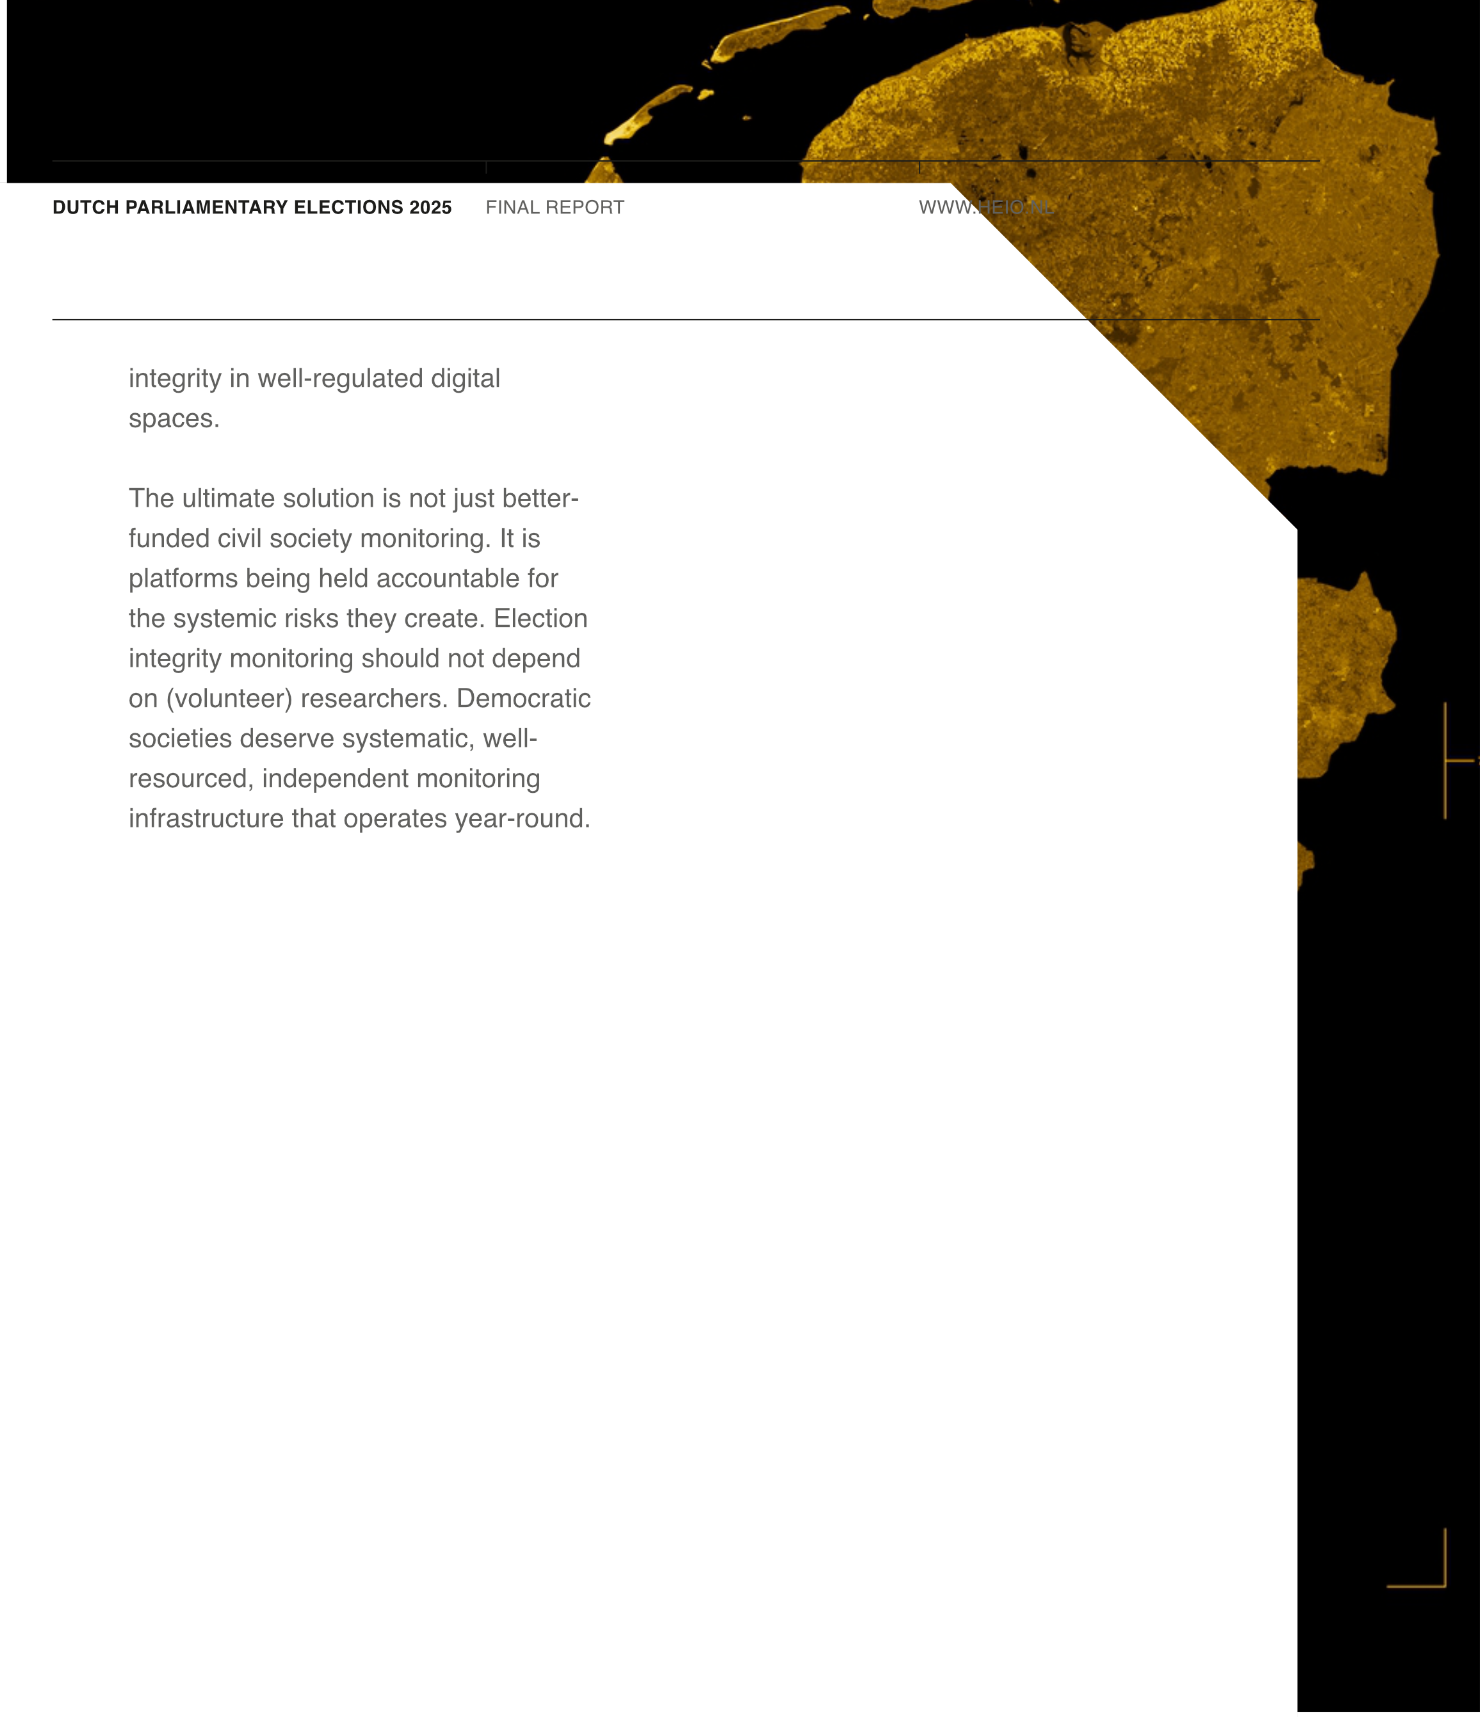

integrity in well-regulated digital spaces.

The ultimate solution is not just better-funded civil society monitoring. It is platforms being held accountable for the systemic risks they create. Election integrity monitoring should not depend on (volunteer) researchers. Democratic societies deserve systematic, well-resourced, independent monitoring infrastructure that operates year-round.

RECOMMENDATIONS

Based on our observations we make recommendations to protect future elections.

ACCOUNTABILITY

» 1 «

SHIFT BURDER OF PROOF TO PLATFORMS

Platforms should be required to demonstrate before, during and after elections through independent certification, that their content moderation systems function effectively. Soft promises must be replaced with verifiable compliance standards.

- Pre-election audits of moderation systems by independent bodies
- Public transparency reports updated weekly during election periods
- Real-time dashboards showing moderation actions and response times

» 2 «

FIX PLATFORM MODERATION SYSTEM

Current moderation is reactive, opaque, and ineffective. Platforms must:

- Increase resources for human moderation in local languages, with transparency on response time
- Publish detailed takedown reasons in standardized, comparable formats

The ephemeral nature of livestream content creates accountability-free zones. Requirements:

- Automatic recording all political livestreams
-
- Real-time moderation with clear, immediate enforcement of Terms of Service violations

» 3 «

PREVENT ALGORITHMIC MANIPULATION

As we noticed new accounts were frequently prominently featured in user feeds, we recommend to:

- Simply prevent very new accounts with political content from appearing at the top of feeds

» 4 «

PREVENT GENAI HATE AND INCITEMENT OF VIOLENCE

CImplement and enforce policies in GenAI tools, such as Sora, that prevent the nonconsensual use of portraits and the generation of content that incites hate or promotes violence against political candidates

TRANSPARENCY

» 5 «

USER ENGAGEMENT TRANSPARENCY

When Elon Musk bought Twitter (now called X) he drastically changed the platform's transparency by turning off that users can see who liked posts. This means that networks/inauthentic behaviour can go undetected and can make posts go viral without the possibility to ascertain who is behind it..

- Reinstate the feature that likes are transparent again.
- Recently X made the geographic location of the user of an account transparent. This feature should be implemented by other platforms as well

» 6 «

POLITICAL ADVERTISING TRANSPARENCY

CThese elections were the least transparent in recent history in terms of campaign financing. Reforms are required:

- Standardized ad libraries with complete content (images, videos, full text)
- Comprehensive coverage of all platforms including "influencer marketing"
- Mandatory disclosure of targeting criteria and spending amountste enforcement of Terms of Service violations

» 7 «

AI CONTENT LABELING

The EU AI Act must ensure that AI-generated content becomes easier to detect, by:

- Mandatory labeling of AI-generated content, especially political content
- Watermarking requirements for AI generation tools, such as Sora
- Video-to-prompt tracing capabilities for synthetic media, in order to reveal the intention of the creator
- Legal liability for platforms that allow unlabeled synthetic political content

REGULATORY REFORM

» 8 «

REAL-TIME ENFORCEMENT

Current retrospective enforcement is inadequate. Platforms also fail to act in the moment, for example when online dynamics contribute to violence at demonstrations. Reforms needed:

- mplement real-time monitoring rather than post-fact analysis
- Set strict timelines: platforms must respond to violations within minutes, not days

» 9 «

REFORM THE RAPID RESPONSE SYSTEM

The EC Rapid Response System needs fundamental redesign:

- Provide full transparency about reporting and outcomes (within privacy constraints)
- Fund civil society participation rather than expecting free labor
- Supplement rather than replace regular user reporting mechanisms

» 10 «

FIX DSA RESEARCH ACCESS

Research access must enable real-time monitoring:

- nsure approval processes are completed before election periods begin
- Enable collaborative research rather than isolating individual researchers, by eliminating unnecessary "clean room" requirements for public data
- Provide standardized data formats and APIs across platforms

» 11 «

ENSURE FUTURE MONITORING

Election integrity monitoring cannot depend on last-minute mobilization:

- Establish dedicated funding for year-round monitoring infrastructure
- Support multi-organization consortia with complementary expertise
- Enable long-term tracking of influence networks and narrative building
- Fund methodology development and tool building between election cycles

Act Now!

Delay is not neutral

It is a choice to leave democratic processes vulnerable to growing manipulation capabilities.

The techniques, networks, and infrastructure identified during the October 2025 elections remain active and are evolving. Municipal elections, provincial elections, and European Parliament elections are on the horizon.

The window for implementing protective measures is narrow.

5

CONTRIBUTORS & THANKS

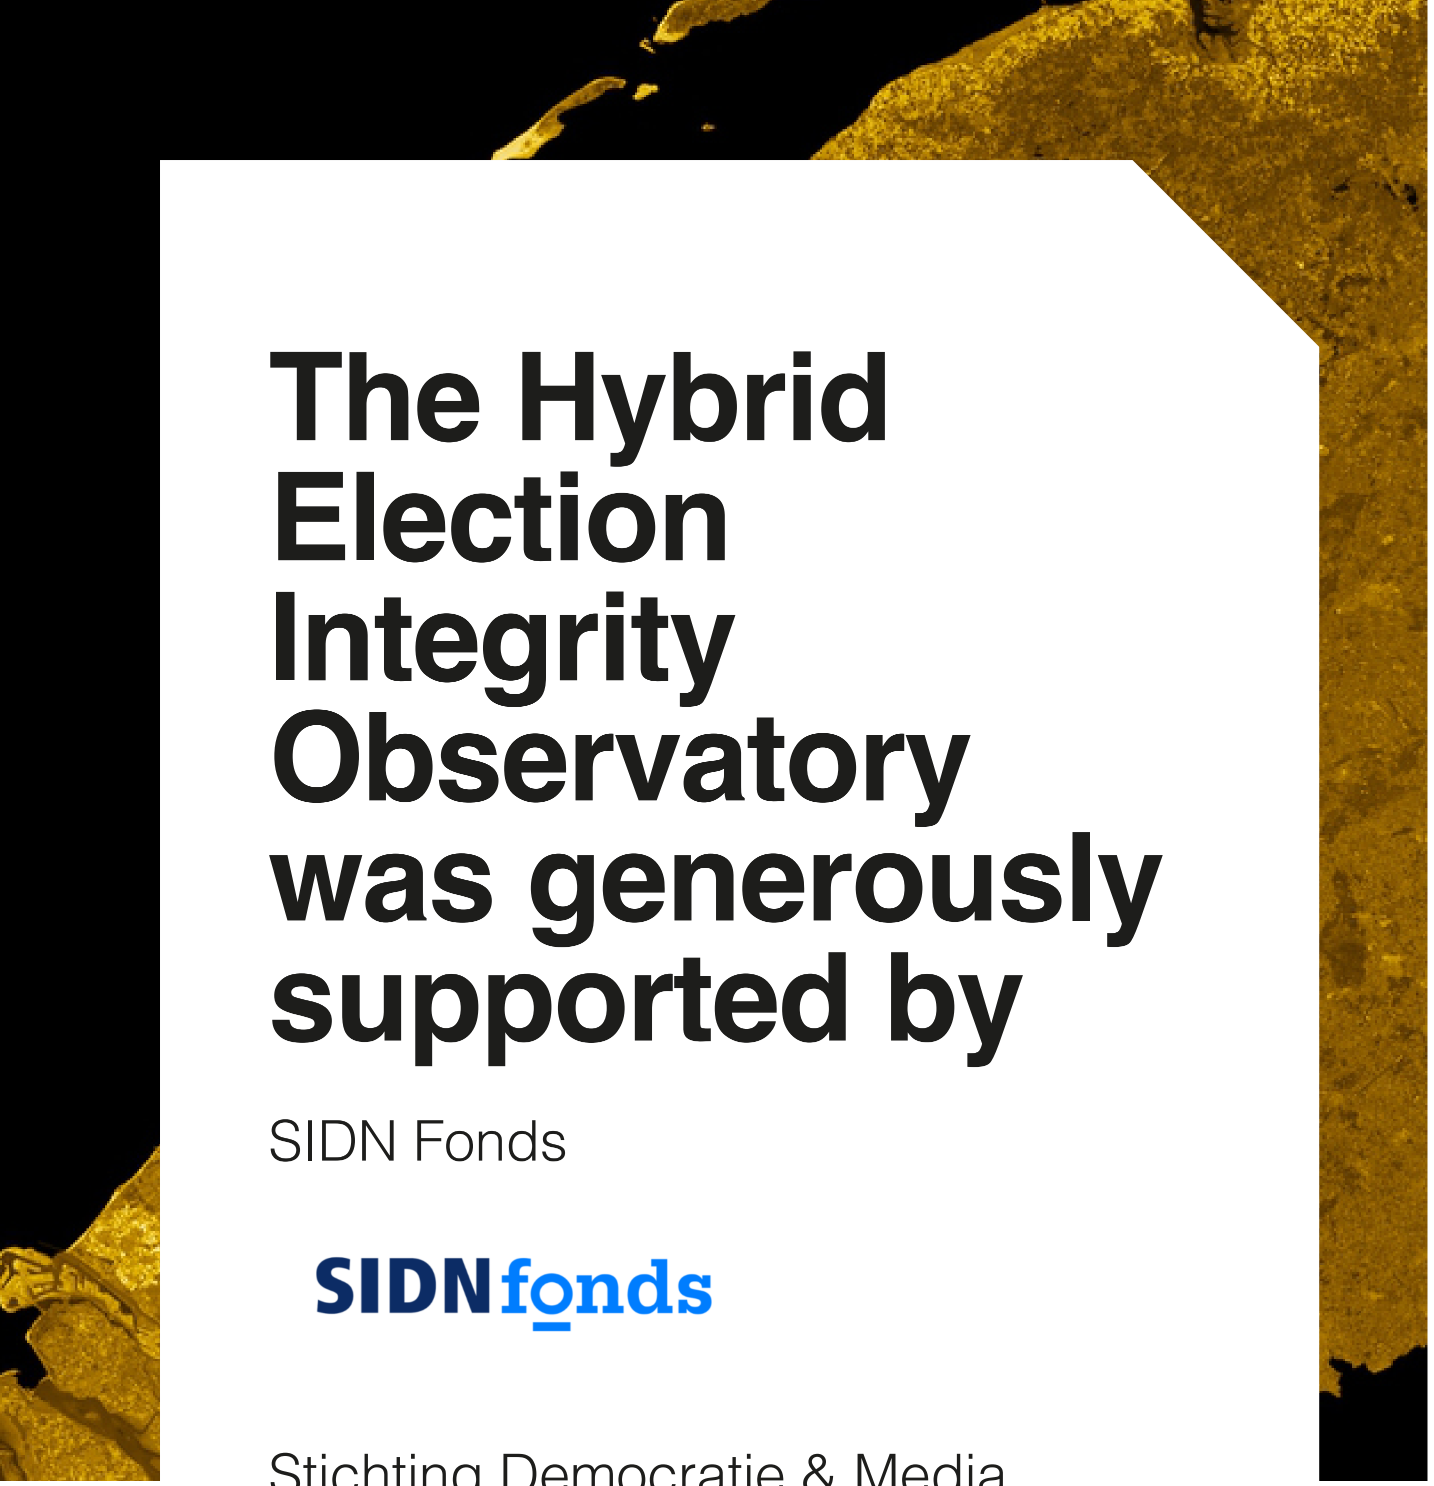

The Hybrid Election Integrity Observatory was generously supported by

SIDN Fonds

SIDN fonds

Stichting Democratie & Media

Stichting Democratie & Media

We would like to thank the citizens who reported leads and shared concerns. Their vigilance and engagement gave us essential insights.

We are also grateful to the investigative journalists at several news outlets, including NRC, RTL Nieuws, and Follow the Money for their efforts to collaborate on, scrutinize and publish our work.

CONTRIBUTORS

The following people contributed to this work:

Post-X Society (PXS)
- Pieter van Boheemen - Director

Trollrensics
- Robert van der Noordaa - Hybrid warfare expert
- Richard Odekerken - Computer scientist
- Jesper Moor – project manager and data analyst

AI Forensics (AIF)
- Natalie Kerby - Researcher
- Miazia Schueler – Researcher
- Katya Viadziorchyk - Head of Fundraising and Advocacy
- Salvatore Romano - Head of Research and Co-Founder

CampAIgnTracker Team (UvA and GESIS)
- Dr. Fabio Votta - Postdoctoral Researcher, AlgoSoc Consortium
- Dr. Simon Kruschinski - Senior Researcher, GESIS - Leibniz Institute for the Social Sciences
- Tobias Kleindamm, Ellen Linnert, Annamijn Beijeman, Haylee Kelsall – Researchers at ASCoR (UvA)
- Daisy Blass, Ella Van Vloten, Doortje Cools, Lynn Louvet, Emily Murphy, Raluca Groza, Isabella van der Zanden - Research assistants

Justice for Prosperity (JfP)
- Jelle Postma – CEO
- Anonymous researchers

DUTCH PARLIAMENTARY ELECTIONS

Hybrid Election Integrity Observatory

PUBLISHED BY POST-X SOCIETY

LANGE HOUTSTRAAT 26
2511CW DEN HAAG

ONLINE
EMAIL: MAIL@POSTXSOCIETY.ORG
WEB: WWW.POSTXSOCIETY.ORG

WWW.HEIO.NL