

Beyond the Label: Generative AI on TikTok and the Limits of Platform Governance in Hungary's 2026 Elections

Executive Summary and Recommendations

The dissemination of AI-generated content on Very Large Online Platforms (VLOPs) is expanding rapidly, particularly in electoral contexts. This shift is reshaping the internet and altering the information environments in which democratic debates take place.

Using TikTok's search function and election-related hashtags, we identified 2,043 posts from 22 accounts that reached more than 50 million views and either posted predominantly AI-generated content or presented themselves as AI-generated personas, in both cases focusing primarily on Hungarian politics and the 2026 elections.

Key Findings:

- Across the 22 monitored accounts, we collected 2,043 posts, 96 per cent of which were published in 2025 and 2026, with a sharp spike in activity around the official start of Hungary's election campaign, on 21 February.
- Nearly three quarters of all collected posts (1,463, 71.61 per cent) were classified as containing AI-generated content and reached substantial audiences. Together, these posts accumulated over 53.6 million "views" (as defined by TikTok), 2.7 million reactions, 1.1 million shares, and 64,000 comments.
- Thirty-seven-and-a-half (37.5) per cent of the AI-generated content identified was not labelled by the user who posted it or by TikTok.
- We reported 199 unlabelled AI posts to TikTok for non-compliance with the platform's Edited Media and AI-Generated Content (AIGC) labelling requirements, and/or for potential violations of policies on synthetic content and disinformation. The company removed 19

videos, one account, and sanctioned two other posts with “visibility restrictions”. The 177 posts that were found not to be in violation are still not labelled, despite our flagging them.

- The large majority of posts (1,295) were Anti-Viktor Orbán, the incumbent prime minister, and Pro-Tisza, the main opposition party, which ultimately won the election. AI-generated content against Orbán also reached much larger audiences per average (42,134 views per post). We cannot ascertain who owned or managed the accounts monitored.** Existing journalistic reporting on AI-generated content in Hungary’s election had largely focused on pro-Orbán use. Our findings, however, show that genAI also played a role in anti-Orbán campaigning.

Recommendations:

To TikTok:

- Strengthen enforcement of the existing AIGC Policy, so that its labelling requirements are consistently applied in practice,** particularly for AI-generated content depicting realistic people or scenes in misleading ways, such as face replacements, voice mimicry, or the fabricated addition or removal of individuals from footage;
- Establish a tiered sanctions framework for labelling violations,** with escalating and clearly defined penalties for repeated failure to label AI-generated realistic-looking people or scenes;
- Revise policies to exclude text captions, hashtag stickers, and post descriptions as acceptable labels.** These formats are not immediately visible to users at the point of contact and prevent the systematic detection and tracking of AI-generated content at scale;
- Provide more information on which posts the platform automatically labels as AI-generated content;**
- Clearly specify what triggers the algorithmic demotion of AI-generated content,** the extent and duration of any reach reduction applied, and how these measures interact with personalised recommendation systems; and
- Support independent research, by granting researchers access to data on AI labels,** including through the APIs provided under Article 40(12) of the EU Digital Services Act.

To the European Commission:

- Under ongoing enforcement processes on generative AI, require VLOPs and Very Large Online Search Engines (VLOSEs) to update community guidelines and terms of service to explicitly address AI-generated political content,** with clear definitions, thresholds, and enforcement pathways, in line with the DSA Electoral Integrity Guidelines;

- ✓ **Require platforms to include AI-generated content labels in researcher APIs under Article 40 (12) of the DSA; and**
- ✓ **Facilitate structured coordination between platforms, civil society, and national authorities to clarify the interaction between the DSA and the EU Artificial Intelligence Act (AI Act), particularly regarding labelling and marking obligations, and ensure coherent implementation across both frameworks.**

Generative AI on TikTok During Hungary's Election Campaign: Key Findings

This report examines the use of generative AI content on TikTok during the 2026 Hungarian parliamentary election campaign. Our monitoring followed a two-step process. We started with a curated list of election-related hashtags (Annex 1), and used TikTok's search function to retrieve relevant content. From there, we identified accounts that were producing generative AI material and, where an account's output was predominantly AI-generated, included it in our dataset. We then traced each account's activity backwards – examining reposts and followed accounts – to uncover further accounts displaying the same patterns. This snowball sampling approach produced an initial set of 40 accounts.

We classified content as AI-generated through systematic manual review, applying established indicators such as visual artefacting, unnatural facial movement, lighting inconsistencies, anatomical distortions, and audio-visual desynchronisation. The annotators discussed “edge cases”, i.e., cases where we were not confident the post contained AI, before assigning a final classification. In most instances, the material consisted of mass-produced AI-generated content, so it was usually easy to classify with confidence.

After reviewing the posts, we narrowed the dataset down to 22 accounts. We included the accounts meeting one or both of the following criteria:

- ✓ **The account posted predominantly AI-generated content; or**
- ✓ **The account profile presented an AI-generated persona, regardless of whether the posted content was predominantly AI-generated**

In both cases, the account also had to meet this condition:

- ✓ **Its most recent content was related to Hungarian politics and/or the 2026 Hungarian parliamentary elections**

We then collected all videos published by these accounts up to 7 April 2026, one week prior to election day.


Account-level analysis

Account creation. More than half of the accounts were created within roughly a year of the election: seven during 2025, and a further five in a concentrated window between 27 January and 24 February 2026, just weeks before the official campaign period began. The remaining accounts were created between 2016 and 2024. That said, even some of these older accounts intensified their output during the campaign period.

Reach. The number of combined followers of the 22 accounts amounted to 121,794, with considerable variation across accounts. The largest account exceeded 60,000 followers, whereas the second-largest had approximately 12,000, indicating a highly skewed distribution. In total, the accounts accumulated more than 4.8 million likes, which were, again, unevenly distributed; the most-liked account alone accounted for around 2.3 million likes.

Activity. We collected a total of 2,043 posts (videos and images) across the 22 monitored accounts. Posting activity was heavily concentrated in the last two months; 1,271 (62.2 per cent) of all posts were published in 2026, and 692 (33.9 per cent) in 2025, with only 80 (3.9 per cent) predating 2025.

The heatmap below illustrates this trend at the account level, revealing that posting activity increased from mid-2025, with several accounts sharply intensifying their activity around the start of the official campaign period, on 21 February 2026.

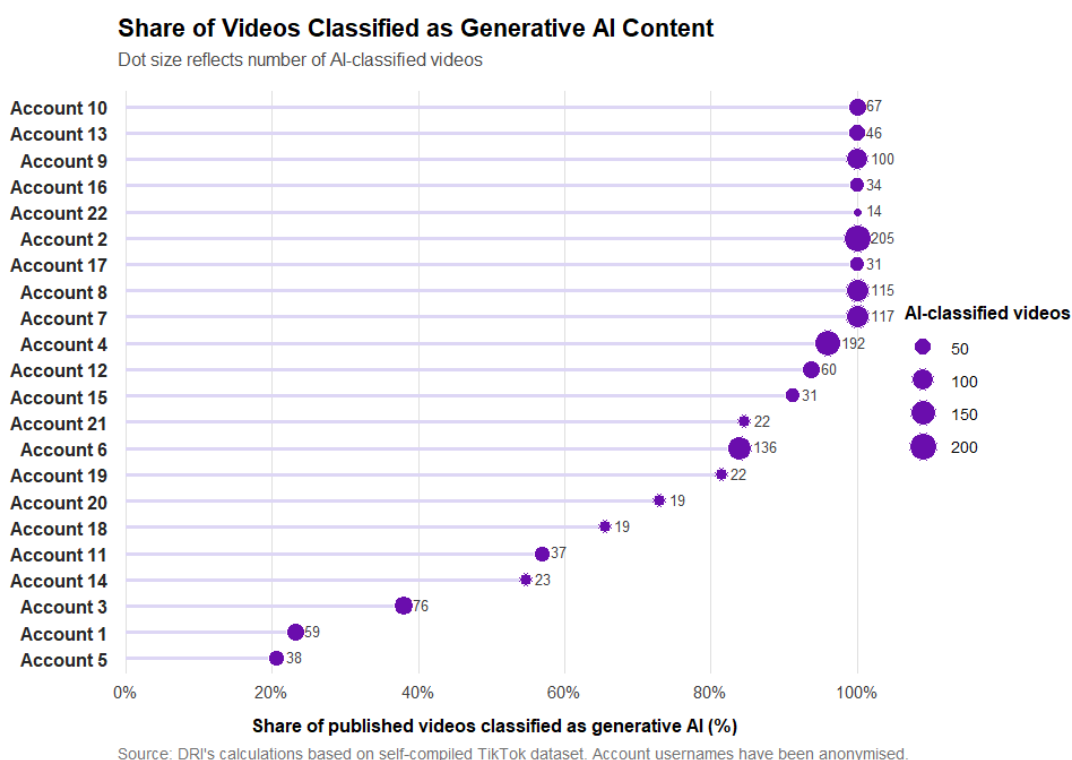


Source: DRI's calculations based on self-compiled TikTok dataset. Account usernames have been anonymised.

Content-level analysis

Of the 2,043 posts collected, we classified **1,463 (71.6 per cent)** as containing **AI-generated content**. This figure should be treated as a conservative lower bound, as the remaining 580 (28.4 per cent) includes posts where AI use could neither be confirmed nor ruled out. The true proportion of AI-generated content across the network could, therefore, be higher.

The graph below shows the share of posts classified as containing AI-generated content, broken down by account. **Fifteen (68.2 per cent)** had **more than 80 per cent of their content classified as AI-generated**. Accounts with a lower share of AI-generated content were, nonetheless, included in the dataset on the basis of the second criterion: their profiles presented an AI-generated persona, making them relevant to the analysis regardless of their content classification.



The 1,463 generative AI videos in the dataset reached an impressive audience. Together, they accumulated over 53.6 million views, 2.7 million reactions, 1.1 million shares, and 64,000 comments.

We classified accounts according to their political orientation through a qualitative analysis of published content, identifying three categories:

Anti-Orbán/pro-Tisza (17): Most accounts produced satirical content targeting Orbán and his party, Fidesz, depicting them in jail, in grotesque scenarios, or as defeated in the upcoming

elections. Some explicitly referenced Tisza through hashtags, campaign slogans, or direct mentions; others did not, yet all consistently framed government change as the central message. Anti-Orbán sentiment and Pro-Tisza alignment proved analytically difficult to disentangle in this context, so both were grouped within a single category.

Pro-Orbán (4): Accounts in this category published messaging critical of Tisza leader Péter Magyar’s campaign and supportive of Viktor Orbán’s continued government.

Anti-Orbán/pro-DK (n = 1): At least one account displayed alignment with the Democratic Coalition (DK, Demokratikus Koalíció), the centre-left opposition party, rather than with Tisza.

Anti-Orbán/Pro-Tisza content was not only the most predominant in our dataset, but also generated the highest overall engagement across reactions, shares, plays, and comments. Satirical accounts consistently outperformed others in reach, suggesting that humour and irony may be more effective at capturing attention than direct party messaging – a finding consistent with existing literature.¹ Whether this reflects genuine political interest or simply entertainment value remains an open question.

The table below presents the average reach per post, disaggregated by the political orientation of the accounts.

Table No. 1.


Average Reach and Engagement per Post by Political Orientation of Accounts

Political Position	No. of Posts	Avg. Reactions	Avg. Shares	Avg. Plays	Avg. Comments
Anti-Orbán/pro-Tisza (17)	1,235	2,165	938	42,134	52
Pro-Orbán (4)	191	136	23	8,330	3
Anti-Orbán/pro-DK (1)	37	19	4	480	1


Source: DRI’s authors own calculations based on a self-compiled TikTok dataset. Calculations are based on the 1,463 posts classified as containing AI-generated content.

Overall, the content posted across these accounts followed patterns already identified in the broader research literature on AI-generated political content. We observed at least four main uses of genAI content:


- AI-generated personas as news anchors.** A recurring format involved AI-generated personas presenting political news directly to the camera, mimicking the conventions of broadcast journalism. In some cases, the persona delivered the entire segment; in others, it served as an introduction before cutting to real footage, interviews, or other non-AI material, blurring the line between synthetic and authentic content;



- Emotionally charged content misrepresenting opposition candidates.** Several accounts produced highly emotive content, designed to distort the perceived positions and character of opposition candidates. A recurring example involved deepfakes of Péter Magyar depicting him engaged in criminal activity, including fabricated scenarios of corruption and of planning to draw Hungary into a war with Ukraine;



- Targeted mockery of specific voter groups.** Several accounts produced content ridiculing identifiable segments of the electorate, rather than political figures directly. A recurring example was the caricature of Fidesz voters as elderly women from rural Hungary depicted receiving bags of potatoes in exchange for their votes or support; and



- Satirical and parodic content targeting candidates and the broader political landscape.**

Most accounts framed Orbán and Fidesz as subjects of mockery and ridicule, though not all of this content advanced a clear political agenda. In several cases, the satire appeared designed less to mobilise support for an alternative than to foster generalised cynicism toward the electoral process as a whole.



Existing journalistic reporting on AI-generated content in Hungary’s election had largely focused on pro-Orbán use. **Our findings, however, reveal a more contested information environment, with generative AI also deployed by supporters of Tisza and Magyar, generating considerable engagement.**

Prior research flagged the use of Orbán’s official accounts to reshare false synthetic content, such as examples implying that opposition candidate Magyar planned to send troops to Ukraine, to divert pension funds to support the war effort, or to introduce forced conscription.²

Carnegie Europe also warned about efforts from Fidesz’ side to invest heavily in semiorganised online activism, including that “networks of AI-generated social media profiles are also producing vast quantities of fabricated and unregulated images and videos promoting government narratives.”³ Similarly, the fact-checking organisation Lakmusz found that the governing party had repeatedly breached Meta’s political advertising ban, including through the use of AI-generated content in paid ads.⁴ Further, Political Capital documented that AI-generated videos were used most intensively by Fidesz and affiliated proxy organisations to amplify the party’s political messaging.⁵

Our findings add nuance to this picture. **Pro-Tisza and anti-Orbán content was not just present on TikTok, but achieved considerable reach, surpassing pro-Orbán content in overall engagement.** This activity did not originate from Tisza’s official accounts or its candidates, but from what appeared to be individual user accounts. It is not possible to ascertain the ownership of these accounts or the intentions behind them.

Research Limitations

Several limitations should be taken into account when interpreting these findings. First, our search methodology (based on election-related hashtags and TikTok’s search function) may have failed to capture relevant content that fell outside our selected search terms.

Second, we deliberately narrowed our sample from an initial pool of 40 accounts to 22, focusing specifically on accounts that published AI-generated content almost exclusively. The 18 excluded accounts also shared AI-generated material, but less predominantly. The amount of AI-generated political content on TikTok during the campaign was, therefore, likely considerably higher than our sample captures.

Finally, AI detection remains an inherent methodological challenge. We classified content as AI-generated or not through manual review, without the support of automated detection tools. This underscores a wider structural problem: without reliable, platform-level labelling of AI-generated content by VLOPs, independent research in this field will remain constrained. Mandatory and machine-readable labelling is not only a transparency requirement, but also a precondition for rigorous electoral integrity research.

AI labelling and TikTok Response



Of the 1,463 posts that we classified as generative AI content, 549 (37.5 per cent) were not labelled, 561 (38.3 per cent) were labelled by the creator using TikTok's voluntary AI content label, and 353 (24.1 per cent) were labelled automatically by the platform.

Following an analysis of TikTok's policies, on 11 April we reported 199 unlabelled AI posts to the platform for non-compliance with its AIGC labelling requirements and potential violations of policies on synthetic content and disinformation. We further recommended that TikTok conduct a systematic compliance review of AIGC labelling practices, and evaluate the flagged content for misleading or deceptive material, with particular attention paid to the electoral context.

On 20 April, TikTok responded. **The company removed 19 videos:** 17 published by one account and two by another; **and took down the account responsible for the 17 removals in its entirety. Two additional posts were sanctioned with visibility restrictions.** The remaining 177 posts were **found not to be in violation.** Notably, at the time of writing, none of the unlabelled posts that were cleared of violations have since been labelled by TikTok, despite our explicit flagging.

Furthermore, by 23 April we had identified that **at least three additional accounts had been removed from the platform entirely, and one had been set to private.** Of the four accounts removed, two were Pro-Orbán and two Anti-Orbán/Pro-Tisza.

How the DSA Governs Generative AI Content on VLOPs

The dissemination of AI-generated content on VLOPs is expanding rapidly, particularly in electoral contexts. This shift is reshaping the political online debate and messaging.

This phenomenon illustrates a dynamic that researchers had anticipated: generative AI dramatically lowers the cost of content production while enabling scalable, high-volume campaign output. Crucially, the persuasive risk of this content does not depend on its perceived authenticity. Even when audiences recognise that content is AI-generated and do not treat it as factually accurate, they may still respond emotionally, engage with it, and redistribute it at scale.

This has a structural consequence that extends beyond individual pieces of content: as AI-generated material becomes more prevalent in users' feeds, it progressively displaces other content, including higher-quality electoral information, or even content from candidates and parties official accounts.

The Digital Services Act has rules to deal with the dissemination of generative AI by VLOPs. In the DSA Electoral Integrity Guidelines, the European Commission suggested the following mitigation measures to reduce systemic risks posed by genAI:

- ✓ **Adapt and enforce terms and conditions** to significantly reduce the reach and impact of GenAI content that spreads disinformation or misinformation about electoral processes. This includes publicly disclosing measures such as labelling, marking, demoting, or removing content;
- ✓ **Clearly label deepfakes** and provide users with an easy way to identify GenAI labels. These labels should persist even after the content is reshared;
- ✓ **Update advertising systems** to allow advertisers to clearly label GenAI content or require such labels for ads; and
- ✓ **Adapt content moderation processes** and algorithmic systems to detect AI-generated or manipulated content, using techniques such as watermarks, metadata, cryptographic methods, logging, or fingerprints.

After an analysis of TikTok's policies on genAI, however, we found some gaps in the way the platform is implementing these mitigation measures.

TikTok requires creators to label AI-generated content in cases where AI is used for realistic-looking people or scenes, for example, when a face is replaced, someone's voice is mimicked, or a person is added or removed in a misleading way.⁶ Yet this research found numerous instances where such labels were absent, suggesting the policy is inconsistently enforced in practice.

Moreover, TikTok’s policies do not specify a sanction ladder for labelling non-compliance in repeated failure to disclose AI-generated content. Without defined sanctions, the labelling regime is unenforceable in practice.

The platform reserves the right to automatically apply its own labels to undisclosed synthetic content, which users cannot remove, but, again, the policies fail to specify what triggers this action or what threshold of AI-generated content warrants a label.

TikTok also allows *the use of texts, hashtag stickers, or post descriptions as acceptable labels*. This approach is concerning as disclosure buried in a post description requires users to actively press “See More”. Most users are unlikely to do this. A label that is not immediately visible when a user sees content functions mostly as a technicality.

Finally, while **TikTok indicates that it may reduce the amplification of misleading content (including AI-generated) through algorithmic measures, they do not specify what triggers demotion**, for how long such measures are applied, or how they interact with personalised recommendation systems. As a result, it is difficult to assess how effective algorithmic demotion is in practice, or how consistently it is applied to AI-generated political content across different users and contexts.

Date: May 2026

About Democracy Reporting International

DRI is an independent organisation dedicated to promoting democracy worldwide. We believe that people are active participants in public life, not subjects of their governments. Our work centres on analysis, reporting, and capacity-building. For this, we are guided by the democratic and human rights obligations enshrined in international law.

Acknowledgements

This brief was written by Democracy Reporting International. This paper is part of the [access://democracy 2026-2028](#) project funded by the Mercator Foundation. Its contents do not necessarily represent the position of the Mercator Foundation.



This publication is available under a Creative Commons Attribution Non-Commercial 4.0 International license.

Julieta Melina Jiménez designed the layout of this publication.

Bibliography

- 1 **Mark Boukes, Xiaotong Chu, M. F. Abdulqadir Noon, Rufei Liu, Theo Araujo & Anne C. Kroon**, [“Comparing user-content interactivity and audience diversity across news and satire: differences in online engagement between satire, regular news and partisan news”](#), *Journal of Information Technology & Politics*, Volume 19, 2022 - Issue 1, 15 June 2021.
- 2 **Reha Kansara & Rita Pálfi**, [“AI videos fuel rhetoric as Orbán bids for four more years in Hungary”](#), BBC, 5 April 2026; AFP, [“AI content proliferates in Hungary ahead of next year’s elections”](#), 14 October 2025.
- 3 **Zsuzsanna Szelényi**, [“New Tricks and AI Tools in Hungary’s High-Stakes Election”](#), Carnegie Europe, 27 January 2026.
- 4 **Lakmusz**, [“The Endgame for Hungary: Disinformation, AI, Prohibited Ads and Russian Interference in the Election Campaign”](#), 2026.
- 5 **Richárd Demény, Péter Krekó, Csaba Molnár, Adél Tilistyák & Konrad Bleyer-Simon**, [“Synthetic influence. Deepfakes and Artificial Intelligence in the Hungarian Election Campaign”](#), Hungarian Media Observatory Against Disinformation (HDMO) and Political Capital, 10 April 2026.
- 6 **TikTok**, [“Community Guidelines: Integrity and Authenticity”](#), last update: 14 August 2025.