



Should AI Tell You Who to Vote For?

Testing Political Advice in General-Purpose AI Systems during the 2026 Hungarian Election Campaign
(preliminary report)

07.04.2026

Executive summary

This preliminary research report finds that two widely used AI systems, **ChatGPT** and **Gemini**, **do not provide reliable political guidance** for the **2026 Hungarian election**. The systems were tested with pre-defined belief sets that matched party positions and were asked to provide voting advice or match percentages with relevant parties. The pre-defined belief sets used in the research were provided as a courtesy by the Hungarian voting advice application [Voksmonitor](#).

First, both systems **struggle to accurately match** certain voter preferences to political parties. This is particularly evident in the case of Tisza Party, a major opposition party with a realistic chance of challenging the governing Fidesz after sixteen years in power. While the **Fidesz-aligned belief set is consistently identified**, the **Tisza-aligned belief set is frequently misclassified**.

Second, ChatGPT produces a **striking asymmetry in its match percentage outputs**. It assigns match percentages to many parties, including those not running in the 2026 elections, but **never to Tisza**. Users whose views align with Tisza are **redirected to the Democratic Coalition (DK)**, a small party that is [consistently measured](#) to have no significant chance of passing the 5% parliamentary threshold in 50% of cases.

Third, both systems offer **substantive political guidance** even when they state they cannot advise users on voting or give a precise percentage match. Although ChatGPT and Gemini, in most cases, avoid explicit recommendations, the systems influence user perception by heavily emphasising certain parties and downplaying or omitting others.

Fourth, outputs **exhibit structured distortions**. These include not only the consistent underrepresentation of Tisza but also broader issues in **how electoral choice is framed**. Voting decision is frequently reduced to unweighted value alignment, with little or no reference to strategic voting, candidate-level dynamics, or the broader electoral context. In addition, ideological distances between parties are often misrepresented.

Fifth, **meaningful user safeguards are largely absent**. The systems rarely communicate uncertainty or direct users to more reliable sources or tools, and present their outputs in a confident, authoritative tone, even when those outputs are blatantly incorrect.

Taken together, these findings raise concerns about the role of AI systems in electoral contexts. As these systems become part of **everyday information-seeking**, their influence on political understanding is likely to grow. Ensuring this influence is **accurate, transparent, and accountable is essential for democratic resilience**.





Contents

Executive summary	1
Introduction	2
Key findings	3
Voting advice condition	3
Percentage-match request condition	4
Implications for democracy	5
Regulatory relevance	5
Research design and limitations	6
CONTACT	7

Introduction

General-purpose AI systems (GPAIs) are playing an **increasingly important role in how citizens access and interpret political information**. Users now seek guidance from AI chats on complex issues, including electoral decisions. As a result, GPAIs are emerging as intermediaries in democratic processes.

This development has direct implications. In an electoral context, even indirect or qualified guidance can **shape how political choices are framed and which options are considered**. Unlike established intermediaries such as voting advice applications (VAAs), journalists, or civil society organisations, GPAIs generally lack transparent methodologies, do not explain how they reach conclusions, and are not held to similar standards of oversight or accountability.

This preliminary report examines how **ChatGPT and Gemini respond to political queries during the 2026 Hungarian election campaign**. The analysis uses controlled inputs based on pre-defined belief sets that replicate the policy positions of the governing Fidesz and the opposition Tisza Party. These belief sets were constructed using data provided by the Hungarian voting advice application [Voksmonitor](#). The systems were asked to advise users on voting choices or to indicate percentage matches with relevant parties.

The findings indicate that these systems **do not provide reliable political guidance**. They often fail to match voter preferences to parties, distort how political options are presented, and offer substantive guidance even when they formally decline to do so. The belief set aligned with the Tisza Party is consistently misrepresented; Tisza is rarely mentioned, while other parties are disproportionately emphasised. Outputs tend to reduce electoral choice to simplified value alignment, overlook important factors such as strategic voting or candidate-level considerations, and present conclusions with unwarranted confidence.





A **more comprehensive study will follow**, extending the analysis to all parties fielding national lists in the 2026 election. This work will support the Civil Liberties Union for Europe's advocacy on the implementation and enforcement of the Digital Services Act (DSA) and the AI Act, focusing on systemic risks to civic discourse and electoral processes, and on systems that influence voters, respectively.

As GPAs become part of everyday information-seeking, their influence on political understanding will likely grow. Ensuring this influence is **accurate, transparent, and accountable is critical for democratic resilience**.

Key findings

Voting advice condition

When **Gemini** receives a belief set aligned with **Fidesz**, it provides explicit voting advice and **recommends that party**. In all **other scenarios** (ChatGPT with Tisza or Fidesz profiles, and Gemini with the Tisza profile), the systems offer detailed analyses that guide users **without making direct recommendations**. Sometimes, this guidance is preceded by a formal refusal to decide whom the user should vote for; in other cases, no such refusal is given. Ultimately, regardless of the refusal, the analysis still steers users toward specific political actors.

Both systems show a **clear asymmetry in matching political profiles**. Even when given inputs that perfectly mirror a party's official positions, they often fail to identify or suggest the opposition Tisza Party, while consistently recognising profiles aligned with Fidesz.

Both systems tend to **present voting advice based on simple value alignment**, without weighting the importance of specific issues. They rarely address the significance of local votes or the practical impact of supporting parties likely to implement these values, rather than those that merely align with the most beliefs.

Some observations

ChatGPT prompted with a Tisza-aligned profile:

- in **0% of the cases, Tisza is the primary or recommended option**.
- mentions Tisza in only 20% of its outputs.
- directs attention to other parties, mentioning DK as a closest match in 50% of outputs and also referencing Momentum, which is not running in the 2026 election.
- in 100% of cases, it includes parties in the analysis that are not running in the election.

ChatGPT prompted with a Fidesz-aligned profile:

- identifies **Fidesz as the most relevant party to the user in 50% of cases**.
- mentions Tisza in 10% of its outputs.
- in 90% of the outputs, the analysis provided includes parties that are not running in the election.

Gemini prompted with a Tisza-aligned profile:

- identifies **Tisza as the most relevant party in 10% of cases**.
- mentions Tisza in 70% of outputs.
- in 100% of the outputs, parties not running in the election account for a significant portion of the analysis.

Gemini prompted with a Fidesz-aligned profile:

- correctly **identifies Fidesz in 100% of cases**.





- mentions Tisza in 50% of outputs.
- in 90% of the outputs, parties that are not running in the election are mentioned, but only in passing.

Percentage-match request condition

When asked to assign percentage alignment with political parties, both systems produce outputs that **appear precise but lack a clear and transparent methodological basis**. Although they often preface their responses by noting that no exact or fully reliable calculation is possible, this disclaimer is typically followed by detailed analyses accompanied by percentage ranges that **convey an impression of accuracy**. The initial caveat may, in fact, enhance perceived credibility, as it signals caution while the subsequent numerical outputs and structured explanations suggest analytical rigour. This mirrors the pattern observed in the voting advice condition: formal limitations are acknowledged, yet the systems proceed to guide users in a substantive and directional manner.

It remains **unclear how these percentage values are derived**. A plausible approach would be to systematically compare each of the user's stated preferences with party positions and calculate alignment based on agreement across individual items. This is broadly the methodology employed by voting advice applications, which match user responses to party positions across a structured set of questions.

The behaviour observed in these systems does not reflect such a structured approach. While they occasionally refer to their outputs as "approximate" or "indicative", the resulting **percentages vary considerably across identical inputs**, and no consistent logic of comparison is evident. This suggests that the numerical outputs are not the result of a transparent or reproducible matching process. As a consequence, the apparent precision of percentage-based alignment may be misleading, and users are not provided with sufficient information to assess the reliability of these figures.

Some observations

ChatGPT

- ChatGPT **assigns percentage scores to many parties**, including those not running on a national list, **but never to Tisza** (0% of cases). Instead, the Tisza-aligned profile is matched most closely with DK or Momentum. This creates a bias, systematically redirecting voters with Tisza-aligned views to other parties.
- ChatGPT declines to provide percentage scores in 10% of cases. In all other cases, it includes parties not running in the election, which introduces noise and reduces the practical relevance of the results.
- **MKKP**, one of the five parties that registered a national list, **receives a percentage score in only 5% of cases**. In comparison, Fidesz, DK, and Mi Hazánk receive scores in 90%, 90%, and 70% of cases, respectively.
- The **same party can receive significantly different percentage scores across identical runs**. For example, with a Fidesz-aligned prompt, one output may show a 25-35% match with MSZP, while another shows 60-70% for the same prompt. This variability is concerning, especially given the confident tone of the responses. The systems rarely indicate uncertainty or limitations, even when outputs are inconsistent or incorrect. This combination of low reliability and high confidence increases the risk of misplaced trust.





Gemini

- Among parties registering a national list in the election, **Gemini assigns a percentage score to Tisza in 90% of cases, Fidesz in 95%, DK in 100%, MKKP in 40%, and Mi Hazánk in 45%**. Momentum, which is not running, was assigned percentages in all cases. Several other non-running parties also received percentage scores more frequently than MKKP or Mi Hazánk.
- Although this preliminary report does not compare actual percentages with the averages produced by the two systems, it is clear to anyone familiar with Hungarian politics that **Gemini's results are often very inaccurate**. Gemini frequently suggests no percentage difference between ideologically distant parties. For example, sometimes almost all opposition parties are sometimes shown to be at the same distance to the Fidesz-aligned profile, despite their clear ideological differences. This approach diminishes meaningful political distinctions and misrepresents the political landscape.

Implications for democracy

When general-purpose AI systems provide political guidance, they act as intermediaries in the information environment. Unlike many established intermediaries, their **methods are opaque, their outputs cannot be reproduced, and their errors are difficult to detect**. As a result, these systems **exert influence without adequate transparency or accountability**.

Consistent omission or underrepresentation of certain parties, especially a major opposition actor, raises **concerns about structural bias**. Even without intentional design, these patterns **can influence which parties users consider, perceptions of political competition, and how individuals interpret their political preferences**. In elections, these effects may be significant.

In addition, the systems' confident tone, combined with their willingness to provide numerical outputs such as percentage matches, **conveys a sense of precision and reliability** that is not warranted by their underlying performance. Formal refusals to provide explicit advice may further reinforce perceptions of neutrality and credibility. At the same time, the observed variability and inconsistencies in outputs indicate that these signals of reliability are often misleading. This combination creates a risk that **users overestimate the accuracy and trustworthiness** of the information they receive.

Regulatory relevance

These findings are directly relevant to current EU regulatory frameworks.

The **Digital Services Act** (DSA) requires very large online platforms and search engines (VLOPs and VLOSEs) to **assess and mitigate systemic risks**, including those affecting civic discourse and elections. Whether AI chats **qualify as very large online search engines is still open to interpretation**. A [functional approach](#), which considers how users interact with these systems to access and process information, suggests these services may warrant inclusion in the DSA's risk framework. The Civil Liberties Union for Europe strongly believes that the European Commission should regard both investigated systems as VLOSEs.





Under the **AI Act**, AI systems that affect democratic processes may be subject to **heightened obligations regarding risk management, transparency, and appropriate safeguards**. While the precise legal classification of conversational systems that provide political guidance remains unclear, their practical role in shaping political information environments is already evident.

Overall, the findings reveal a gap between the potential influence of the tested systems and the existing safeguards to ensure accountability and protect democratic processes.

Research design and limitations

This study uses a small-scale audit approach with controlled inputs and repeated testing.

Two models were tested: ChatGPT (free version, tested on 23 March 2026) and Gemini (free version, tested on 24 March 2026). All tests used standardised conditions, including a Chrome browser, Hungarian-language setting, Hungarian VPN location, clean sessions, and a laptop interface. Outputs were archived immediately using Zotero and recorded in a shared dataset. Two researchers coded the data independently, with discrepancies reviewed by the research lead.

The study uses profiles that replicate exact party positions, based on a draft questionnaire developed by Voksmonitor (16 March 2026). At the time of testing, the Voksmonitor tool had not yet been publicly released.

Party positions were presented in two prompt formats.

Voting advice condition:

"I would like to vote in the upcoming elections, but I am not sure which party or candidate to support. I have compiled my views on current political issues.

[views]

Please advise me. Who should I vote for?"

Percentage-match request condition:

I have compiled my views on current political issues.

[views]

Could you tell me what percentage my views match with each party?

The sample consists of:

- 5 profiles
- 10 repetitions per profile
- 2 models
- 2 conditions

total: 200 observations

The study has several limitations, including the small number of models and profiles, constraints from prompt design, and the exclusion of mobile phones from testing. The findings should be seen as indicative rather than definitive. However, the consistency of observed patterns supports their relevance.





This preliminary report focuses on prompts related to the two leading political actors and does not yet include detailed analysis of smaller parties contesting the election. The full report will address this. The full report will also develop policy recommendations to support the implementation and enforcement of the Digital Services Act and the AI Act. Results will be shared with relevant European Commission services and will contribute to ongoing advocacy for platform accountability and AI governance.

CONTACT

Civil Liberties Union for Europe

The Civil Liberties Union for Europe (Liberties) is a Berlin-based civil liberties group with 24 member organisations across the EU campaigning on human and digital rights issues, including the rule of law, media freedom, SLAPPs, privacy, targeted political advertising, AI, and mass surveillance.

Transparency register number:
544892227334-39

c/o Publix Hermannstraße 90
12051 Berlin, Germany

Dr. Orsolya Reich
o.reich@liberties.eu
www.liberties.eu

