Policy brief

European Partnership for Democracy

AI and electoral processes: what's new with the European Democracy Shield?

10 September 2026

Introduction

The **European Democracy Shield** (EUDS), the EU's umbrella strategy for democracy support, was published in November 2025. It rests on three pillars: safeguarding the integrity of the information space; strengthening democratic institutions, including free and fair elections, and free and independent media; and boosting societal resilience and citizens' engagement.

As such, the EUDS also addresses issues at the **intersection of digital policy and democracy**, offering an overarching strategy that ties together several existing initiatives in this space. In recent years, a number of regulations have been introduced to tackle specific digital challenges to democracy. For instance: the **Digital Services Act** (DSA) for platform regulation, the **General Data Protection Regulation** (GDPR) for data protection, the **Political Ads Regulation** (TTPA) for political ads online, and the **Artificial Intelligence Act** (AI Act) for Artificial Intelligence. What has been missing is a structure linking these pieces together specifically from the standpoint of protecting democracy, since initial attempts such as the European Democracy Action Plan (EDAP) have only scratched the surface. Rather than adding new digital rules, the EUDS fills this gap by providing an **overarching strategy.**

The EUDS includes few new elements, but one exception is the forthcoming **Guidance on the use of AI in electoral processes**. It will complement existing regulation by clarifying what fair, transparent, human-centred, and responsible use of AI in elections looks like. The Guidance will be prepared by the European Commission with support from Member States, ECNE, and the European AI Office, and will encourage voluntary commitments from political parties and other relevant actors on the responsible use of AI, while promoting the sharing of best practices. As such, it offers a valuable opportunity to tie together digital rules and

electoral processes for a more effective protection of free and fair elections from the impact of AI systems.

This brief provides further background on the current state of regulation around Artificial Intelligence and democracy, particularly on electoral processes, and examines the role the Guidance could play in complementing and strengthening the existing framework.

Who's afraid of AI in elections: risks and opportunities

It has been widely documented that **AI has an impact on democratic processes** across several dimensions, including civic engagement and online discourse, transparency and accountability of public administration, free and fair elections, and the broader exercise of fundamental rights such as freedom of expression and the right to privacy. In this context, we focus specifically on the impact of AI on **electoral processes**.

AI systems can affect electoral processes - positively or negatively - because they are deployed in different ways before, during, and after elections. Several mappings of AI applications in this space have been put forward, sometimes organised around how citizens experience AI through phases such as production, interaction, communication, and campaigning, each associated with specific use cases. Drawing on these mappings, the most prominent uses of AI in electoral contexts include, among others:

- **Voter data analysis and predictive analytics**, used to perform microtargeting of voters based on inferred preferences and behaviour (the Cambridge Analytica case being the most known example);
- **Ad delivery and microtargeting**, with AI used to tailor political advertising to specific audience segments based on behavioural or demographic data;
- **AI chatbots**, either general-purpose or specialised chatbots deployed by electoral authorities or campaigns to answer voter queries and guide participation; and/or specifically aimed at helping citizens navigate registration, polling locations, and procedural requirements;
- **Generative- and general-purpose AI**, used to produce campaign material, sometimes including synthetic text, images, audio or video (with associated risks of disinformation and deepfakes);
- **Agentic AI** is a recent development that allows accounts to be operated autonomously with an AI-generated persona. They publish AI generated content exclusively, including in the context of electoral campaigns (in the 2026 Hungarian elections, for instance), but not only. What's generated here is not just the content, but the whole account.
- **General-purpose and specialised recommender systems**, which shape what political content users see on social media and other platforms, influencing exposure to information and opinions;
- **Content moderation**, deployed to counter biased or harmful content and to identify disinformation.

- **Ballot processing and voter list verification**, applying AI to count or process ballots and to maintain and verify electoral rolls;
- **Cybersecurity**, with AI used to detect and respond to attacks on the IT infrastructure underpinning elections;

These technologies offer **genuine opportunities**. AI can strengthen electoral administration and integrity, for example by improving the **accuracy and efficiency** of ballot counting or voter list maintenance and by helping detect irregularities. It can support the **detection of disinformation** itself, through automated tools that flag suspicious content or coordinated inauthentic behaviour. AI-based tools can also foster civic engagement and participation, by making information about candidates, procedures, and polling logistics **more accessible**, including for voters with disabilities or language barriers. Finally, AI can provide valuable **policy and legislative support**, helping electoral management bodies and lawmakers analyse trends, anticipate risks, and design more effective regulatory responses.

At the same time, these use cases also rise to **a set of recurring risks**. AI can accelerate the distribution of harmful content, including **AI-generated disinformation** and deepfakes. It can enable more sophisticated election interference, through **microtargeting and behavioural manipulation** techniques that go beyond traditional campaigning by tailoring messages to individuals. The large-scale collection and processing of voter data needed for many of these applications also raises significant personal **data protection concerns**. AI systems trained on historical or unrepresentative data can also **amplify existing biases**, for instance in how content is moderated or how voters are targeted.

More broadly, the complexity of these systems often results in a **lack of transparency, explainability, and accountability**, making it difficult for regulators, candidates, or voters to understand how decisions or content were generated, which can contribute to a wider distortion of public opinion.

We are not yet doomed: a portfolio of policy solutions to address main challenges

Beyond mapping the risks and opportunities associated with AI in elections, work has also been done on identifying the **policy measures that could most effectively address them**. The measures mentioned most often include:

- **Labelling, watermarking and disclosure of AI-generated content**: requirements for AI-generated content to be labelled as such, or for AI chatbots to disclose that users are talking to a machine rather than a human. This obligation typically falls on providers of AI technologies and can be done on the output (visual watermark on an AI generated image) or on the model level (leaving marks that can be identified by detection tools, better for writing or audio). Labelling doesn't always prevent the underlying harm (an AI-generated nudified image of a candidate still causes damage

even when labelled), so this measure needs to be combined with others below. Some of these requirements already exist in the EU's AI Act for specific categories of AI systems.

- **Platform detection and monitoring obligations**: requirements for online platforms, where AI-generated content spreads, to identify and label it accordingly. Similar caveats apply as above, with the added complication that platforms are generally not obliged under existing frameworks to conduct "general monitoring" of hosted content, only targeted checks or risk-prevention measures (see DSA obligations below).
- **Accuracy and reliability requirements, including on training data**: requirements for AI providers to train models on high-quality, up-to-date, non-biased datasets, so that outputs are correspondingly accurate and reliable. Additional requirements on output quality could be introduced specifically for electoral contexts. Training data requirements already appear in the AI Act for high-risk systems, while reliability requirements aren't explicitly spelled out but could fall under the mandatory risk management systems required of high-risk AI providers.
- **Transparency and auditability of recommender systems**: obligations for platforms to disclose the main criteria their algorithms use to rank content in users' feeds, increasing accountability around the spread of AI-generated content and the microtargeting of political ads through manipulative techniques (part of the DSA and, partially, the AI Act).
- **Content moderation quality**: transparent standards for how platforms moderate hosted content, increasing accountability around the spread of AI-generated content (part of the DSA — though this addresses transparency of the moderation process rather than the content itself).
- **Rules on AI-driven advertising**: ensuring that transparency rules for online political advertising also apply to advertising embedded in AI chatbot responses, correctly labelled as such. This is implicit in the political ads regulation and the DSA, but it's unclear how far it extends to AI chatbots, which aren't technically online platforms.
- **Protections against candidate impersonation**: certain AI applications could be banned outright if they offer no value beyond the potential to mislead or damage voters, with candidate impersonation as a leading example.
- **Incident response and institutional coordination mechanisms**: both platforms and AI providers should maintain internal mechanisms to detect, respond to, and coordinate on incidents as they arise.

Some of these measures are already, at least partially, **reflected in the current EU framework**, which we will further explore in the next sections.

What has been done so far: existing EU regulation on AI

None of these issues are new: they have been widely documented, and initial efforts have been put in place to address them, both in EU regulation and beyond. It is important to note, however, that until the Democracy Shield and the upcoming Guidance on AI in electoral processes, there had been **no encompassing EU approach to AI in elections**. This is not to say that the topic wasn't addressed, but individual issues were split across broader pieces of legislation and hidden under the banner of product safety or transparency, rather than being integrated in a specific framework dedicated to democracy.

The two most relevant frameworks to have considered AI and elections so far are the **Artificial Intelligence Act** and the **Digital Services Act**.

Artificial Intelligence Act

The Artificial Intelligence Act (AI Act) is the EU Regulation introducing **rules for AI systems based on their level of risk**. While it has no specific purpose of protecting election integrity from the use of AI systems, it can nonetheless serve as a tool to help ensure free and fair elections, by shielding them from the potential negative impact of certain AI applications.

As a horizontal framework, the AI Act has the potential to affect a wide range of issues linked to fundamental rights, including the right to vote. The Regulation itself states, both in its recitals and in Article 1, that among its objectives are the protection of democracy and the rule of law, as well as the protection of the fundamental rights enshrined in the Charter of Fundamental Rights of the European Union.

Several sections of the AI Act refer, explicitly or implicitly, to **AI systems with the potential to affect election integrity**, and to the safeguards attached to them depending on their risk level. In particular:

- **Prohibited AI systems**, those that cannot be deployed on the EU market at all. This includes for example social scoring, manipulative techniques directed towards vulnerable categories and categorisation of people based on political opinions. It remains unclear whether any of the use cases outlined above would fall into this category.
- **High-risk AI systems**, which must comply with specific obligations such as conducting risk assessments and implementing risk management measures. This category explicitly includes AI systems "intended to be used to influence elections," which may capture some of the use cases described above. A caveat here is that while some systems do indeed influence elections, it is difficult to prove that this is their intention.

- **Limited-risk AI systems**, which are subject to specific transparency obligations, such as labelling and watermarking. Most of the systems flagged above as potential threats to elections, including general-purpose AI and chatbots, fall within this category.

The AI Act - which fully entered into application on 2 August 2026 (with the exception of the provisions for high-risk AI systems) - therefore contains many important provisions to ensure that AI systems are safe to use in the context of elections. The **effectiveness of these provisions**, however, will fully depend on their interpretation and how many use cases will be captured by the new rules. The concrete implementation of measures such as risk management and watermarking will also make a difference on the impact of the legislation.

Digital Services Act

The Digital Services Act (DSA) is the EU Regulation setting obligations for online platforms, with rules focused mainly on transparency. While it has no specific focus on AI, several of its provisions nonetheless affect AI applications, particularly the dissemination of AI-generated content on social media. For instance, the DSA sets transparency requirements for **recommender systems** (Article 27 and 38), and requires Very Large Online Platforms and Search Engines to conduct **risk assessments** and implement mitigation measures across several areas, including **civic discourse and electoral processes** (Article 34). This would include risks related to the dissemination of AI generated content online.

Building on this, the Commission has also released **Guidelines on electoral processes**, setting out best practices for mitigation measures, including practices around AI-generated content, such as labelling and watermarking.

Overall, the DSA has a lot of potential to ensure electoral integrity by mitigating the risks posed by AI, but its scope is **limited to the output of AI systems**, such as the spread of AI generated content online and doesn't address how AI systems are designed and built from the ground up.

What's next? Towards comprehensive Guidelines on AI in electoral processes

Based on the work that has been done so far, there are specific interventions that could be put forward both to make existing rules more effective in practice and to consider dedicated new policy measures for categories that remain only partially or indirectly addressed. In particular, the following regulatory tools should be taken into account in the AI Guidelines in order to include some of these measures:

1. **Consider specific AI systems as high risk and put forward risk management measures according to the AI Act:** AI systems such as ad delivery and microtargeting, specialised

recommender systems, and AI chatbots in certain cases should be considered as high risk under Annex III 8b AI Act. This would oblige AI providers to put forward risk management measures such as a risk assessment; limiting certain features when it comes to electoral content; accuracy requirements; high quality database of training data; and transparency on the functioning of AI systems. The high risk designation does not condemn AI systems in their entirety, but serves to mitigate both intended and unintended risks that could result from them.

2. **Implement transparency measures under the AI Act and DSA**: the AI Act provides transparency obligations around labelling and disclosure that should be respected for limited and high risk AI systems (chatbots and general purpose AI). These are also further specified in the GPAI Code of Conduct and Transparency Code of Conduct, both integrated into the AI Act. The DSA Guidelines on electoral processes also similarly mandate transparency obligations on AI generated content that is spread on social media platforms, in particular around watermarking. While transparency itself is unlikely to mitigate risks, it is crucial for policy and decision makers to be able to create suitable policies and practices as AI systems continue to develop.

3. **Consider some systems used in the context of elections as prohibited under the AI Act**: systems that could feed into the category are for example AI systems for voter data analysis and predictive analytics to perform microtargeting (when they use subliminal manipulative techniques); and general-purpose AI systems, such as chatbots & virtual assistants, if they use subliminal manipulative techniques in sensitive contexts (under Articles 5.1a, 5.1b and 5.1g).

4. **Address risks associated with AI chatbots**: recent developments showed that AI chatbots are increasingly used to obtain information about elections, political parties and related topics - and that they tend to be biased and contain factual inaccuracies. For these reasons, it would be important to have a more comprehensive strategy going beyond the transparency measures outlined above. For example, AI chatbots used to obtain information about elections should be part of the high-risk category in the AI Act, so that related risk management measures can be put in place such as accuracy standards, redirecting to official channels and even refusing answers in certain cases.

5. **Alternative governance models for AI**: For an even more comprehensive approach it would also be worthwhile to consider promoting alternative governance models for AI applications. AI technologies are controlled by a handful of top-down companies who do not embed considerations of public interest into their AI products. In the context of the EU's efforts towards digital sovereignty it would be important to promote more bottom-up, participatory, community based AI solutions (including open source) so that a plurality of interests could be included from the get go. This would be an important step in ensuring that AI could really be deployed for the common good.

Conclusions

AI's impact on elections has so far been addressed as part of broader sets of rules, such as the AI Act and the DSA, rather than through rules designed specifically for elections. This leaves specific **gaps open**, as outlined above, for example on whether key use cases qualify as high-risk or prohibited under the AI Act and what are the most effective mitigation measures under the DSA.

The Democracy Shield and the upcoming Guidance on AI in electoral processes are an important step towards a genuinely encompassing framework. But to be meaningful, this step must **systematically address the core issues identified** above in coordination with existing regulation.

More broadly, this requires **moving past siloed tech policy**. Industrial policy and product safety can no longer be treated as separate from their impact on democracy: safeguarding **elections needs to be built into AI governance from the outset**, not added as an afterthought.