



ETHICAL DILEMMAS OF AI-GENERATED POLITICAL CONTENT IN DIGITAL POLITICAL COMMUNICATION

1 Tianxinran, 2 Dr. Mohd Haizra Hashim

^{1 2} Faculty of Business and Communication, INTI International University, Nilai, Negeri Sembilan,
Malaysia

Email: I26041058@student.newinti.edu.my, mhaizra.hashim@newinti.edu.my

Abstract

Generative AI revolutionizes digital political communication, empowering everyday citizens, political players and media outlets to create highly realistic political content, both text and multimedia. As these tools have the potential to enhance communication effectiveness, creative expression and public political participation, there are also significant ethical concerns, such as high-volume misinformation, intentional manipulation of audiences, low accountability and the reproduction of social bias and a growing loss of public political trust. Upon this basis, four well-established ethical approaches—Utilitarianism, Virtue Ethics, Rights Based Ethics and the Potter Box Model will be used to dissect these fundamental ethical issues. Particular focus is paid to contexts in Malaysia and Southeast Asia, where social media platforms are the key and most important public space for political debate. The authors of this study contend that responsible AI-mediated political communication needs clear disclosure of AI content, meaningful human oversight, strengthening platform accountability, systematic fact checking measures and enhanced media AI literacy among citizens. Rather than just the technology, ethical responsibilities should be distributed among a variety of political actors, AI technology developers, digital platforms, regulatory bodies, and general users.

IndexTerms-artificial intelligence, political communication, deepfakes, media ethics, misinformation, digital media

I. INTRODUCTION

Generative AI has reduced technical hurdles, enabling users with no technical expertise to access advanced synthetic media. In the field of political communication practice, political parties and candidates use these to reach out to the public in multiple languages, to design campaign materials and to engage with the audience through interactive tools. At the same time the technical capacity can be used to create false political narratives and audiovisual manipulated content (Battista, 2024; Chowdhury, 2024).

Political communication has particular social sensitivity since citizens need political information to make their opinions and informed democratic judgements. However, it is difficult for common viewers to effectively identify whether content is real or AI-generated by politicians (Barari et al., 2025; Groh et al., 2024). The moral debate is thus not the one concerning the use of AI technology in political communication in general. The focus, rather, is on the deployment of generative AI systems and the maintenance of factual truth, human autonomy, procedural fairness and public trust. While AI-generated translations, satirical parodies, and educational simulations are legitimate applications when clearly labelled, serious ethical issues arise when audiences are not informed that the content they are consuming is artificially generated or manipulated (Romero & Moreno, 2024).

Generative AI significantly lowers the economic barriers to creating election campaign content, while current political communication rules and regulations are not very well prepared to face new synthetic media threats (Łabuz & Nehring, 2024). This paper addresses key ethical challenges raised by AI-generated political content, critically analyzes via four lenses of ethical theories, and suggests context-specific responsible practice practices suitable for the Malaysian and other Southeast Asian digital political landscapes.

Research Questions RQ1: What sets of core ethical dilemmas emerge from AI-generated political content in digital political communication contexts? RQ2: In what ways can Utilitarianism, Virtue Ethics, Rights-Based Ethics and the Potter Box Model be used to assess these ethical dilemmata? RQ3: What are the responsible communication measures that suit the digital political context of Malaysia and Southeast Asia?

II. LITERATURE AND CONCEPTUAL BACKGROUND

A. AI-Generated Political Content

AI-generated political content covers synthetic photographs, voice-cloned audio, AI-authored political texts, virtual political figures and algorithm-tailored political advertisements. Generative AI makes synthetic media production more accessible and less costly, enabling private citizens, grassroots political groups, and social media influencers to produce and distribute political content on a wider scale, without necessarily relying on the expertise of professional campaign teams or traditional media outlets (Battista, 2024).

While AI opens up new avenues for engaging linguistically and demographically diverse audiences, it also makes it easier to produce and distribute misleading content. Political actors are creating favourable images of themselves through the use of manipulated media, even if it means not actually staging events in order to do so – Chowdhury (2024) calls this “political softfakes”. Softfakes are different from high profile deepfakes. Softfakes can be something as simple as image manipulation or recontextualised clips, while deepfakes use powerful AI to create realistic human faces and voices. While the technical complexity of both depends on the individual, both can deceive an audience.

Therefore, it is important to recall that ethical evaluation of media should be applied to the entire spectrum of altered media and not just to malicious deepfake productions. There are a number of interdependent factors that influence whether the content will have a real-life impact: the intent of the creators, the extent of media manipulation, the transparency of content labeling, the susceptibility of the audience, and the predictable consequences in the real world (Hameleers, 2024).

B. Political Deepfakes and Information Disorder

Deepfakes are media that have been manipulated or based on AI synthesis that show real-world people saying or doing things that they didn't actually do. Ordinary viewers can easily miss the detection of many political deepfake clips, as shown in controlled experimental research (Groh et al., 2024). Despite this, deepfakes don't necessarily make for more convincing fake information than traditional text-based fake information. Their effect depends on the credibility of the source and the plausibility of the message and on the prior attitudes of the audience members (Hameleers, 2024).

An additional, but very dangerous side effect is the “liar's dividend”. Once synthetic media becomes well known, it may be easy for political actors to discredit real audio or video recordings by falsely attributing them to synthetic media. This practice compromises the accountability of institutions and makes it more difficult for the public to confirm political misconduct in the real world (Chesney & Citron, 2019).

C. Political Trust and Democratic Communication

Healthy democratic political-communication fundamentally depends on public trust in information authenticity. Exposure to synthetic political videos can lead to higher uncertainty and lower overall trust of news products (Vaccari & Chadwick, 2020). But even if audiences do not fully trust the content they encounter, there can be substantial societal damage if audiences remain in a state of uncertainty, which undermines the common source of information that democracy depends on (Hameleers et al. , 2024).

Altogether, the literature indicates that the issue of AI-generated political content poses an ethical challenge not only on the audience's belief with regard to a specific manipulated political message. It also includes uncertainty and source evaluation and/or trust of the overall information environment. Thus, short-term impacts on audience judgement need to be taken into account, as well as longer-term implications for the political communication.

III. THEORETICAL FRAMEWORK

A. Utilitarianism

Utilitarianism calls for evaluating human behavior according to the net effects. In the context of media ethics, the communication practice has to create more aggregate social benefits than aggregate harms, in order to be considered ethically acceptable.

AI-powered political communication brings concrete benefits such as cost reductions in production, multilingual communication and greater civic engagement. However, the potential harms are not to be overlooked: misinformation, reputational damage of individuals, political polarization, voter manipulation, violation of privacy, and erosion of political trust in the long run. A practical perspective is that short term electoral or campaign involvement is not worth the certain long term impact on society. This framework is designed to ensure that communicators think beyond the campaign results and consider risks at the community level.

B. Virtue Ethics

Virtue Ethics is not just an evaluation of results or adherence to rules, but rather one that focuses on the moral character of the actors. The most important virtues that are applicable to political communication practice include honesty, procedural fairness, responsibility, prudence, transparency and respect for audience agency.

A political communicator doesn't have to make an outright false claim about the facts, but can use AI-generated visual tools to create a false impression overall. Virtue analysis thus seeks to determine whether communicators behave with integrity and take proper responsibility for the effects of what they publish downstream. Responsibility doesn't just lie with politicians. Synthetic media technologies developed and hosted by AI developers and digital media platform intermediaries also come with moral responsibilities. Even having the ability to create synthetic media technically cannot free organisations of their moral responsibility (Pawelec, 2025).

C. Rights-Based Ethics

Based on the concept of fundamental human rights—access to reliable information, personal privacy, human dignity, and freedom of thought and of meaningful civic participation—Rights Based Ethics evaluates communicative acts.

Audience's ability to make informed political decisions is compromised when they can't tell if an image is manipulated with AI. Additionally, the unauthorized creation of a person's likeness or voice may lead to issues of personal dignity, privacy, and individual autonomy (Chesney & Citron, 2019; Romero & Moreno, 2024). In this context, citizens need to be considered as rights-holders who are democratic actors and not just passive objects of political argumentation.

D. Potter Box Model

The Potter Box Model is a four-stage ethical decision-making model that can be used in complex value conflict scenarios:

Situation: Generative-AI systems create or substantially revise political-communication materials.

Core competing values: truth, fairness, individual autonomy, privacy and public.

Principles: Guiding ethical principles: minimise unnecessary harm, reject deliberate deception, respect individual human rights and promote responsible public communication.

Loyalties: Primary loyalty is to the public good, not to the ballot, audience numbers or the bottom line. For instance, in a political deepfake scenario, it will be synthetic content being disseminated to shape political actions. These values are electoral effectiveness, truthfulness, autonomy and public welfare. The ethical principles call for communicators to avoid knowingly causing harm or deception and the question of loyalty raises the question of whether communicators should put the goals of the campaign first or the interests of the public to make informed choices. This indicates that the Potter Box can be applied to detect when conflicts arise between a political interest and a role as a public interest.

The competing stakeholder interests are also brought into focus by this model. Ethically questionable communication decisions cannot be based on their effectiveness as a message or the success of the campaign.

IV. CRITICAL ETHICAL ANALYSIS

A. Truth, Accuracy, and Misinformation

Political videos created by AI can be used to distort the actual words of politicians and mislead large audiences, even when the politicians issue official rebuttals. Corrected fact checks are not necessarily seen by everyone who initially consumed the manipulated short video, particularly on widely shared platforms such as short videos (Barari et al. , 2025).

There is still a strong need to distinguish between deliberate disinformation and creative and clearly labelled applications of AI (such as parody or educational simulation). The main moral risk is misleading presentation and intentional generation of false, public appearances. Politically motivated content that is manipulated to fit with audiences' existing beliefs is likely to receive additional credibility, which will boost its actual impact in the real world (Hameleers et al. 2024).

Case: 2024 New Hampshire AI Generated Biden robocall: Thousands of robocalls using an AI voice that impersonated Biden were deployed in the United States 2024 New Hampshire primary election to encourage people not to vote in the election (Swenson & Weissert, 2024). Consideration of utility led to the conclusion that the electoral confusion and threats to electoral integrity were greater than the narrowly partisan advantage. The actors in this robocall were being dishonest and un-civic, from a virtue ethics point of view. According to the principle of the rights-based ethics, the right to access to a genuine election-related information on the part of the voters has been violated. In the context of the Potter Box Model, narrow electoral tactical goals were taken over with regard for public interest.

B. Transparency and Accountability

When political media content has been created or significantly manipulated using AI technology, audiences need to be informed about it. If not disclosed, ordinary viewers will naturally believe that the photos, sound and video are true real-world events. This makes it difficult for audiences to make a critical assessment and to assign responsibilities when harm arises.

There are also many different actors involved in the production and dissemination of synthetic political content: those who create the content, those who commission a political campaign, providers of the AI technology, those who publish or distribute content and end users who re-share. The sharing of ethical responsibility should be spread through the entire synthetic media ecosystem rather than focusing on individual end users (Pawelec, 2025). Meaningful disclosure needs to be easily understandable for general audiences and cannot be hidden inside technical metadata or fine-print captions.

C. Manipulation, Persuasion, and Conflict of Interest

In a democratic political media, persuasion is a normal and legitimate part of communication. Candidates are entitled to persuade voters to vote for their policy positions. But, for persuasion to be morally acceptable, the basic respect for the autonomy of the audience must be adhered to.

The boundary between legitimate persuasion and harmful manipulation becomes blurred with AI-enhanced political content. Persuasion clearly argues something so that an audience is able to evaluate the evidence and come to their own conclusions. In contrast, manipulation plays on a group's emotions, psychological weaknesses, and lack of context and information to deceive citizens and to reduce their ability to make an informed judgment (Chowdhury, 2024). Using AI-powered workflows, creating personalised, emotionally charged manipulative content is easier at scale.

Political actors often have conflicts of interest: the need to have a successful campaign may be at odds with the need to provide truthful information to their constituents. If the primary objective in the electoral process is winning, then content generated by AI but not disclosed moves from persuasion to manipulation. In regulatory responses, distinctions should be drawn between a prohibition of all persuasive political speech and a prohibition on deliberate deception, harmful impersonation and foreseeable harm to democratic participation.

D. Political Trust and Democratic Legitimacy

Repeated exposure to synthetic political content can have significant, systemic consequences for society as a whole. Citizens slowly lose confidence in the public sources of information from which they draw their information (Vaccari & Chadwick, 2020).

As shown by Hameleers et al. (2024), even technically implausible content can be created in the form of deepfake productions for targeted politicians, which have a delegitimising effect. When coupled with the liar's dividend effect, synthetic media forms a serious second layer of threats to the functioning of democracy: politicians can challenge authentic recordings that incriminate them by claiming they were AI-generated, thus undermining institutional accountability. When citizens can no longer trust what they hear as genuine political conversation and what is fabricated, it becomes more difficult to participate meaningfully in politics.

E. Bias and Fairness

Patterns in AI system outputs are the product of training data, which are likely to contain social, cultural and political biases of real-world society. Political communication by AI can thus perpetuate negative social stereotypes and generate unfair representations of social minorities. Bias can occur at various points in the production process, such as data selection, generation pipelines, and algorithmic recommendation systems.

There must be more than just automated systems in place for political communication to be fair. AI-generated content should never be assumed to be neutral or unbiased. Political communicators need to evaluate if synthetic content is exacerbating harmful stereotypes and/or undermining specific communities (Battista, 2024).

F. Privacy, Surveillance, and Political Micro-Targeting

AI in political communication raises significant privacy concerns when linked to extensive data analysis of users and segmenting audiences based on their behavior. Political micro-targeting is the ability to tailor political communication through segmented groups of citizens, but at the same time it raises serious concerns about mass personal data exploitation and citizens' awareness.

From a Rights-Based Ethics viewpoint, citizens should not be reduced to mere behavioural data inputs for campaign use. Political communication should meet basic privacy requirements and personal autonomy (Chesney & Citron, 2019; Romero & Moreno, 2024). The risks are exacerbated when hyper customised messaging is tailored to prey on group vulnerabilities. It is good campaign practice to minimise personal data collection where this is not necessary.

V. AI-GENERATED POLITICAL CONTENT IN THE MALAYSIAN AND ASIAN CONTEXT

Malaysia offers a good non-western case study. Its multilingual, multicultural society relies heavily on social-media platforms as the main arena for political discussion. The rapid dissemination of short-form visual media on WhatsApp groups, TikTok and Facebook makes it easy for synthetic political media to spread like wildfire.

Sabri (2024) notes that these issues of AI-generated deepfakes and the dissemination of misinformation are not limited to the western world but are also genuine threats to Malaysian politicians and the integrity of national information. In recent Malaysian state election cycles, several image and audio clips were shared that were manipulated using unauthorised AI technologies, targeting politicians in the country. Some manipulated content to alter the meaning of politicians' speeches or to create false visual scenarios to harm reputations or to enhance partisan narratives (Sabri, 2024). Comparative studies across Asia also reveal that the public's sense of threat from deepfakes is significantly influenced by the degree of institutional trust in the region rather than technological attributes alone (Neyazi et al. , 2024). Technical one size fits all solutions will not suffice. Localised responses should also take into consideration local political sensitivities, multilingual moderation issues, the dynamics of the video platform and the varying levels of digital literacy among different community groups in Malaysia.

VI. RESPONSIBLE COMMUNICATION PRACTICES

Clear AI Disclosure: Clear, prominent and easy to understand labelling of campaign content with prominent AI elements that have been generated or substantially modified by AI. Particularly for synthetic voices, face renderings and re-enactment style political videos there are requirements for labelling. Transparent disclosure helps maintain audience sovereignty to make a critical judgement of political speech.

Human Oversight: AI tools should never be used as a substitute for human ethical considerations. Political groups must have in place formal procedures for human review of political content created by AI. Factual accuracy, foreseeable harm, fairness, transparency and overall ethical compliance are all tested by human reviewers (Pawelec, 2025).

Digital platforms should have clear, public policies on synthetic political media; have available content reporting processes and respond in proportion to deceptive harmful content. Platform governance should be adequate to reconcile objectives of harm prevention with freedom of expression and safeguard legitimate satire, commentary and educational content.

Privacy Protection: Political campaigns should collect as little personal data as is necessary. Political targeting using AI should be done in a transparent way and should not be based on sensitive personal attributes or psychological weaknesses.

Flows: Journalists and fact checking organisations should have better verification workflows in terms of suspicious political audiovisual material (Walker et al., 2024). Purely technical AI-detection tools are not enough. The verification should include automated scanning, expert human review and cross source context-based checking. Reliance on regular viewers isn't going to adequately resolve the problem of manipulation once content has already gone viral.

Media and AI Literacy: Public education should aim to foster more general critical thinking skills, such as verifying where the contents are coming from, what the producer's intent is and what evidence is provided, rather than just teaching the public to identify obvious visual artefacts created by deep fakes. Most important, literacy education does not abdicate responsibility for identifying synthetic media to citizens, because humans' ability to identify sophisticated synthetic media is well documented (Groh et al., 2024).

VII. STAKEHOLDER RESPONSIBILITIES

Political actors and campaign organisations: Make it a priority to communicate truthfully, transparently reveal the use of AI, avoid deceptive strategies, and take responsibility for the content generated. The electoral victory should never replace the right of the citizens to make their informed political decisions.

Digital platforms: Develop clear synthetic media policies, enhance content reporting systems and carry out targeted risk assessments at key election times. Recommendation algorithms that affect the reach of content, have equivalent ethical responsibilities.

AI developers: create technical measures to prevent malicious impersonation, plan potential political misuse and incorporate ethical aspects in the design of the technologies (Pawelec, 2025).

Governments and regulators: adopt proportional rules to prevent electoral integrity being undermined while at the same time not limiting legitimate free expression. Regulation focus needs to be on deliberate deception, harmful impersonation, privacy breaches, and real threats to democratic engagement.

Media organisations and journalists: Check content in doubt before making it political, make people aware of the dangers of synthetic media.

Educators and higher education institutions: provide structured media literacy training and AI media literacy courses for critical examination of digital political content.

General audiences: Be wary of sharing unverified political material that might be emotionally charged without thinking first and confirm key facts before spreading them on. These ethical dilemmas cannot be solved by a single stakeholder alone, and therefore the cooperation of multi-parties in governance is needed.

VIII. DISCUSSION

AI technology is inherently neutral. There are differences in the way the four frameworks address the ethical issue. Utilitarianism takes into account the consequences of political communication produced by AI, and virtue ethics takes into account the communicators' honesty and responsibility. When using rights-based ethics, the emphasis is on autonomy and dignity/privacy while with Potter Box the emphasis is on conflicts between values, principles and stakeholder loyalties. Therefore, taken in combination, these frameworks give a wider, more comprehensive moral evaluation than any one of the frameworks alone. It is not technology that creates core ethical conflicts but human decisions and responsibilities. Utilitarianism focuses on trying to maximize the benefits of communication and minimize known social costs. Virtue Ethics focuses on the communicators' moral character such as honesty, integrity and prudence. Rights-Based Ethics emphasizes the importance of citizens' rights to accurate political information, to privacy, to dignity and to autonomous civic decision making. The Potter Box Model provides a framework to deal with clash of values and multiple stakeholder preferences.

The existing academic literature validates that real-world consequences of deepfakes are situational, meaning that the plausibility of the message, the preexisting attitudes and institutional trust of the audience are significant factors Barari et al. (2025); Hameleers et al. (2024). There is no solution to this social issue that relies solely on technical detection tools since the technology of generative AI continues to advance quickly, and bad actors are constantly innovating their deceptive content.

Asian context studies, particularly in Malaysia, show that effective governance responses should be based on local cultural and institutional contexts (Sabri, 2024; Neyazi et al., 2024). A complete blanket ban would suppress legitimate AI applications such as translation, accessibility aids and political satire. Full deregulation, on the other hand, allows for extensive deceit and manipulation. So the best practical solution is responsible transparency: allowing for use of AI-assisted communication where legitimate, and implementing clear and specific protections against extreme or harmful manipulation or impersonation and intentional public deception. Freedom of expression is a prerequisite for democratic political communication, as is also social responsibility.

IX. CONCLUSION

While AI can help streamline political processes, enhance creativity and expand reach in political communication, it also poses significant ethical challenges, including: misinformation, manipulative persuasion, transparency issues, accountability lapses, privacy worries and loss of public political trust.

Through the application of Utilitarianism, Virtue Ethics, Rights-Based Ethics and the Potter Box Model, it has been revealed that political communication should be implemented ethically and has to be characterised by honesty, shared responsibility, respect of citizens rights, harm reduction and prioritisation of the public interest. The impact of synthetic political media can't be boiled down solely to technological impacts; impacts on real world outcomes are influenced by source credibility, situational contexts, audience interpretation, platform distribution behaviours, and general institutional trust.

A focus on the Malaysian and Southeast Asian context reveals that AI political ethics discourse needs to capture the social reality in the context. Practical measures such as compulsory disclosure of AI-generated content, formal oversight by humans, greater platform liability, privacy protection policies and better fact checking systems, as well as media literacy about AI, are considered meaningful.

This study is based entirely on secondary-sources literature. Future empirical studies may look into the public perception in Malaysia with regards to the political content created by AI, as well as how the levels of political trust and media literacy affect the ordinary citizens' assessments of synthetic political media.

REFERENCES

- Barari, S., Lucas, C., & Munger, K. (2025). Political deepfakes are as credible as other fake media and (sometimes) real media. *The Journal of Politics*, 87(2), 510–526. <https://doi.org/10.1086/732990>
- Battista, D. (2024). Political communication in the age of artificial intelligence: An overview of deepfakes and their implications. *Society Register*, 8(2), 7–24. <https://pressto.amu.edu.pl/index.php/sr/article/view/42182>
- Chesney, R. M., & Citron, D. K. (2019). Deepfakes: A looming challenge for privacy, democracy, and national security. *California Law Review*, 107(6), 1753–1820. <https://cyberlaw.stanford.edu/publications/deep-fakes-looming-challenge-privacy-democracy-and-national-security/>
- Chowdhury, R. (2024). AI-fuelled election campaigns are here—where are the rules? *Nature*, 628, 237. <https://doi.org/10.1038/d41586-024-00995-9>
- Groh, M., Sankaranarayanan, A., Singh, N., Kim, D. Y., Lippman, A., & Picard, R. (2024). Human detection of political speech deepfakes across transcripts, audio, and video. *Nature Communications* 15, 7629. <https://doi.org/10.1038/s41467-024-51998-z>
- Hameleers, M. (2024). Cheap versus deep manipulation: The effects of cheapfakes versus deepfakes in a political setting. *International Journal of Public Opinion Research*, 36(1), edae004. <https://doi.org/10.1093/ijpor/edae004>
- Hameleers, M., van der Meer, T. G. L. A., & Dobber, T. (2024). Distorting the truth versus blatant lies: The effects of different degrees of deception in domestic and foreign political deepfakes. *Computers in Human*

Behavior, 152, 108096. <https://doi.org/10.1016/j.chb.2023.108096>

- Łabuz, M., & Nehring, C. (2024). On the way to deepfake democracy? Deepfakes in election campaigns in 2023. *European Political Science*, 23, 454–473. <https://doi.org/10.1057/s41304-024-00482-9>
- Neyazi, T. A., Nadaf, A. H., Tan, K. E., & Schroeder, R. (2024). Does trust in government moderate the perception towards deepfakes? Comparative perspectives from Asia on the risks of AI and misinformation for democracy. *Government Information Quarterly*, 41(4), 101980. <https://doi.org/10.1016/j.giq.2024.101980>
- Pawelec, M. (2025). Decent deepfakes? Professional deepfake developers' ethical considerations and their governance potential. *AI and Ethics*, 5, 2641–2666. <https://doi.org/10.1007/s43681-024-00542-2>
- Romero & Moreno, F. (2024). Generative AI and deepfakes: A human rights approach to tackling harmful content. *International Review of Law, Computers & Technology*, 38(3), 297–326. <https://doi.org/10.1080/13600869.2024.2324540>
- Sabri, F. A. A. (2024). Narrative review on detonation of AI-generated deep fakes and misinformation: A threat to politicians in Malaysia. *e-Journal of Media and Society*, 7(2), 71–88. <https://ir.uitm.edu.my/id/eprint/97612/>
- Swenson, A., & Weissert, W. (2024, January 22). New Hampshire investigating fake Biden robocall meant to discourage voters ahead of primary. AP News. <https://apnews.com/article/new-hampshire-primary-biden-ai-deepfake-robocall-f3469ceb6dd613079092287994663db5>
- Vaccari, C., & Chadwick, A. (2020). Deepfakes and disinformation: Exploring the impact of synthetic political video on deception, uncertainty, and trust in news. *Social Media + Society*, 6(1). <https://doi.org/10.1177/2056305120903408>
- Walker, C. P., Schiff, D. S., & Schiff, K. J. (2024). Merging AI incidents research with political misinformation research: Introducing the Political Deepfakes Incidents Database. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(21), 23053–23058. <https://doi.org/10.1609/aaai.v38i21.30349>

