

Merging AI Incidents Research with Political Misinformation Research: Introducing the Political Deepfakes Incidents Database

Christina P. Walker, Daniel S. Schiff, Kaylyn Jackson Schiff

Purdue University
walke667@purdue.edu

Abstract

This article presents the Political Deepfakes Incidents Database (PDID), a collection of politically-salient deepfakes, encompassing synthetically-created videos, images, and less-sophisticated ‘cheapfakes.’ The project is driven by the rise of generative AI in politics, ongoing policy efforts to address harms, and the need to connect AI incidents and political communication research. The database contains political deepfake content, metadata, and researcher-coded descriptors drawn from political science, public policy, communication, and misinformation studies. It aims to help reveal the prevalence, trends, and impact of political deepfakes, such as those featuring major political figures or events. The PDID can benefit policymakers, researchers, journalists, fact-checkers, and the public by providing insights into deepfake usage, aiding in regulation, enabling in-depth analyses, supporting fact-checking and trust-building efforts, and raising awareness of political deepfakes. It is suitable for research and application on media effects, political discourse, AI ethics, technology governance, media literacy, and countermeasures.

Background

Recent Advances and Political Implications

Deepfake generation capabilities have evolved significantly in recent years, with the technology now capable of creating hyper-realistic content with minimal expertise (Barari, Lucas, and Munger 2021). This includes not only synthetic text, but also audio, images, and videos, making it increasingly difficult for individuals to reliably distinguish between real and fake content (Groh et al. 2022; Rana et al. 2022). While deepfakes and other synthetic media, including simpler ‘cheapfakes,’ serve diverse purposes from entertainment and education to fraud, a class of particularly concerning uses are politically-oriented deepfakes. Recent instances include the use of deepfakes to critique political figures such as Donald Trump, Joe Biden, Emmanuel Macron, Nancy Pelosi, Vladimir Putin, and Volodymyr Zelenskyy.¹ Deepfakes have been leveraged to influence voter turnout (Christopher 2020), accuse politicians of sexual scandals

(Cook 2019), and even influence geopolitical events like the Russian invasion of Ukraine (Simonite 2022).

Additionally, political campaigns need to be aware of the potential use of deepfakes. Candidates in the 2024 U.S. presidential primary races have already created campaign ads involving deepfakes (Nehamas 2023). This is especially problematic as broadcast networks cannot refuse to run a campaign ad, as The Communications Act of 1934 says they have “no power of censorship” over any “legally qualified candidate” for office (Montanaro 2022). In response, the Federal Election Commission voted unanimously to pursue evaluating the rules around AI use in campaigns (Klar 2023).

Consequences can be far-reaching. For example, a fabricated video depicting a Pentagon bombing caused a dip in the U.S. stock market (Sorkin et al. 2023). In another instance, mere accusation of a deepfake helped facilitate a coup in Gabon (Cahlan 2020). Despite efforts by leading artificial intelligence (AI) providers to curb the creation of politically harmful content, researchers have demonstrated the ease with which mainstream tools can be bypassed. For example, Midjourney prompts like “evil politicians grinning, sad children, Comet pizza shop” and “George Floyd realistic robbing a Wal-Mart” (Alba 2023), referencing a political conspiracy theory and a contentious high-profile police incident in the U.S., respectively, can be used to generate fake images potentially capable of stoking real political divisions. Consequently, scholars have urged that deepfakes could extend the harms of mis- and disinformation. This includes influencing elections, furthering democratic backsliding, undermining trust in media and government institutions, fostering pervasive uncertainty, inciting intolerance and violence, undermining public health, providing new tools for disinformation campaigns, harming psychological security, and empowering authoritarian governments (Sharma et al. 2023; Goldstein et al. 2023; Vaccari and Chadwick 2020; Schiff, Schiff, and Bueno 2023).

Policy Responses

International, national, and subnational entities, along with private sector actors, researchers, civil society groups, and standards organizations have already started to develop policies and regulations to address these concerns. Responses range from deepfake detection methods and watermarking to policy taskforces, platform regulation, media literacy ef-

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹These deepfakes are included in the PDID, which can be found at: <http://bit.ly/pdid>.

forts, and outright bans (Rana et al. 2022; Partnership on AI 2023). The European Union, for instance, has been advancing efforts to regulate the use of AI systems including deepfakes through risk assessments and transparency requirements (European Commission 2021). China has enacted regulations of deepfakes with criminal penalties (Hao 2023). The U.S. has also considered a deepfake task force and secured a commitment from leading AI companies to watermark AI content, furthered by a 2023 Executive Order (Stokols and Egan 2023), and Congress has proposed the DEEPFAKES Accountability Act. Similarly, states and cities are also exploring legislative measures to curb the misuse of deepfakes, with nine U.S. states enacting regulation as of 2023 in lieu of federal action (Williams 2023).

Although some deepfake databases have been developed, such as the FakeCatcher database, AI-Powered Video Archive, and the Deepfake Detection Challenge Dataset, these are largely aimed at advancing technical detection methods and not comprehensive across the spectrum of politically and publicly salient deep- and cheapfakes, making systematic analysis challenging. The growing legislative concern along with the increasing sophistication of deepfakes underscores the need for a comprehensive incidents database, which aims to provide the means for valuable insights into the prevalence and impact of political deepfakes.

We thus introduce the Political Deepfakes Incidents Database (PDID). The fundamental goal of the database is to provide a publicly available resource to advance research, practice, and governance efforts surrounding deepfakes. Journalists, fact-checking organizations, media literacy educators, and members of the public can utilize the database to evaluate the veracity of specific instances or identify broader trends of concern. In addition, the PDID can assist in understanding the effectiveness of watermarking and detection techniques, media literacy efforts, and other policy initiatives on deepfake dissemination and impact.

Creating a Political Deepfakes Database

Disciplinary Gaps in Deepfake Research

With the increased usage of generative AI in politics, there is a corresponding need to advance research into political deepfakes. Much research to date has focused on whether deepfakes are convincing to viewers (Barari, Lucas, and Munger 2021), whether they are persuasive in terms of changing beliefs or attitudes (Schiff, Schiff, and Bueno 2023; Wittenberg et al. 2021), and whether they can be detected (Groh et al. 2022). Yet many of these studies take place in idealized laboratory studies using research-created content, undermining our capacity to understand real-world behavior and effects. Thus, observational data are needed to assess the prevalence and impact of ‘real’ deepfakes with heightened political salience that *have gained public attention*.

Related but broader concerns about the potential harms of AI have led to the initiation of a research agenda around AI incidents. Examples include the AI Incidents Database which now features more than 500 incidents as of 2023 (McGregor 2021), the Awful AI Database (Dao et al. 2022), and

the AI Vulnerability Database². Reporting certain kinds of incidents may soon become a requirement by international regulators, though potentially limited to incidents caused by products and services of companies, rather than adversarial actors or private individuals. Indeed, one challenge identified by this research is the existence of many more reported ‘issues’ where harm is not as easily detectable (McGregor, Paeth, and Lam 2022).

Political misinformation is one such case, where serious consequences are possible, such as social upheaval, political instability, collapses in institutional trust, and ascendance of unaccountable or authoritarian leaders. However, the pathway to these harms is difficult to identify causally; individuals may be exposed to large amounts of content over time including messages and counter messages, making it difficult to identify single deepfakes as unitary sources of harm. Information-seeking behavior and attitudinal change may differ across demographic subgroups, such as across regions or groups with different levels of digital literacy. Harms may thus result from pervasive inputs acting through complex causal chains, with feedback effects and heterogeneous effects rendering analysis difficult. Nevertheless, given the growing importance of this research as misinformation increases and democracy backslides, determining strategies for identifying and mitigating deepfake-related harms is critical. The growing toolkit of misinformation researchers can be useful here to support work being done on AI incidents collection and analysis, currently primarily by technical researchers or AI generalists. For example, scholars in political science, communications, and related fields have developed and tested theories and models to evaluate issues like:

1. How actors or institutions are incentivized to promote misinformation/disinformation
2. How misinformation is shared, promoted, and accessed through various channels
3. How different types and characteristics of policy frames or narrative strategies can strategically shape views of issues, politicians, or social groups
4. How individual and group political psychology, including in-group and out-group behavior, influence the efficacy of misinformation
5. How the timing of messaging and counter-messaging leads to persistence effects or decay effects as contrasting messages compete or are reinforced
6. How observational, quasi-experimental, and experimental evidence can be mustered to test the validity of theories, including studies of the efficacy of policy changes, mitigation efforts, and choices around AI system design

To enable effective research on political deepfakes and mitigations then, information about the prevalence and nature of deepfakes must be curated and aligned with theoretical elements that a broader set of researchers can draw on to advance practical knowledge and associated action.

²<https://avidml.org/>, accessed 01/25/2024

Purpose of the Political Deepfakes Incidents Database (PDID)

To foster these goals, this article introduces the Political Deepfakes Incidents Database (PDID). The goal of this database is to create a comprehensive and interoperable collection of political deepfakes and cheapfakes, beginning with videos and images³ primarily from the U.S.⁴, to facilitate research on their usage in politics. The PDID is motivated by and aims to draw on the informational needs of misinformation researchers and practitioners—in media studies, political communication, artificial intelligence, public policy, and political science, amongst others—to enable the study of the prevalence, nature, and impact of political misinformation. By incorporating deepfakes, their content and context, real-world connections, and related aspects of politics and policy coded by researchers, the database aims to provide a valuable resource for understanding the prevalence, trends, and potential impact of political deepfakes. Critically, the PDID enables empirical studies helping to fill one of the key gaps in current research: the lack of causal evidence linking the use of political deepfakes to real-world outcomes. The database thus holds potential for various stakeholders, including policymakers, researchers, journalists, fact-checkers, and the public.

Policy leaders can benefit from this database by gaining insights into the scope and scale of political deepfake manipulation. Understanding how deepfakes are used in both international and domestic politics can help policymakers formulate targeted regulations and measures to address the challenges posed by this technology while safeguarding democratic processes and national security. Scholars and analysts can utilize this database to conduct in-depth analyses and investigations into the patterns and dissemination of political deepfakes. By studying the content, context, and impact of these manipulated materials, researchers can enhance our theoretical understanding of the societal, political, and psychological implications of deepfakes.

For example, in the field of communication, the PDID can provide insights into how deepfakes are crafted and disseminated and how they influence public discourse. Information technology practitioners can use the database as a resource for developing and testing new algorithms for detecting deepfakes. Political scientists can investigate the role of deepfakes in shaping political narratives, influencing elections, and destabilizing democratic processes. Fact-checking organizations and journalists can use the PDID to cross-reference the authenticity of content and to report on trends. Finally, the PDID can help educate the public about the existence and potential dangers of political deepfakes.

³We omit audio-only content initially in order to set a realistic scope for the first phase of the PDID. Individuals tend to believe visual content due to a strong “realism heuristic” (Vaccari and Chadwick 2020), so we focus on the misinformation that arguably has the largest potential impact on trust in the informational environment. We may incorporate modalities like audio in the future.

⁴We also initially limit our scope to the U.S., with intentions to scale the database gradually, likely relying on the on-the-ground expertise of partners in other countries and regions.

Methodology

Data Collection and Codebook Development

Data collection began in June 2023 and is ongoing, while the PDID currently includes deepfakes incidents dating back to 2017. The database is hosted on the platform AirTable,⁵ which the research team used to manually code the variables described below for each deepfake incident. Data are primarily collected from English language social media posts and popular news websites given our focus on deepfakes that have reached the public eye. To capture associated data, we begin with these news sources reporting on prominent deepfake incidents and use a snowball sampling technique, looking at linked articles and social media posts.

To create the structure for the PDID, we began by collecting an initial sample of prominent political deepfakes, which we define as deepfakes featuring or commenting on prominent political figures, both individuals and groups, as well as politically important social issues and events. We then iteratively developed a codebook to guide data collection and coding for the PDID, for coders, and later analysis and interpretation of the database. We determined the relevant categories based on the data we collected for the initial sample, and we identified various theoretically important indicators based on recent political science and communication research and external feedback. To improve interoperability and enhance rigor towards greater impact, we consulted existing and forthcoming coding taxonomies developed by the AI Incidents Database, Center for Security and Emerging Technology, and AI Vulnerabilities Database.⁶

The criteria that we used to develop the codebook are as follows. First, the variables needed to parse features of deepfakes with sufficient variation and granularity to enable meaningful study (e.g., the harm type depicted). Second, the coding choices for the incidents needed to be reliable and replicable (e.g., an original source could be identified with sufficient confidence based on information in the deepfake, key social media post, or news article). Finally, certain variables were identified because they held potential theoretical importance for misinformation researchers (e.g., the framing of a deepfake). We thus divided the codebook into categories including metadata, verification data, entity descriptors, and social media and sharing information, along with variables surrounding context and content, real-world connections, and theoretical codes related to politics and policy.

Database Structure and Variables

As noted, we incorporated variables from the AI Incidents Database and CSET taxonomy that we thought were both relevant and analytically tractable, removed inapplicable items, and added breadth and depth where appropriate. For example, while CSET’s taxonomy has a special focus on identifiable and often physical harms for substantive and methodological reasons, the PDID codebook was designed to capture theoretically relevant aspects of political deepfakes and thus was modified to capture less ‘tangible’ political and social harms.

⁵The database is available at: <http://bit.ly/pdid>.

⁶We thank these teams for sharing their insights on the PDID.

Descriptors	URL, File, Media format, Summary
Social Media	Views, Likes, Comments, Shares
Source	Name, Entity type, Location
Sharer	Name, Occupation, No. of followers
Target(s)	Name, Entity type, Response (if any)
Verification	Presentation as real or fake, Visual markings of false content, Watermark/label, External sources
Context & Content	Text from post, Context from post and comments, Harm type depicted (e.g., violence or discrimination), Communication goal (e.g., satire or reputational harm)
Real-world Connections	Context (e.g., scandal, election, or war), Reported real-world harms (e.g., evidence of actual violence or political interference)
Politics & Policy	Policy sector, Framing, Narrative

Table 1: Overview of Variables in PDID

Table 1 provides a simplified overview of the variables included in the PDID, and the online Appendix⁷ provides full variable lists and definitions, which we expect to evolve over time. Below, we provide detail on the motivation and construction of some key variables of interest.

Metadata: For metadata, we capture information such as the fake itself (e.g., deepfake or cheapfake, video or image), whether the incident has been verified as fake or authentic, the level of social media engagement (e.g., number of likes), and when available, information on the type, name, and location of the target and sharer entities.

Verification: We collect information regarding whether the images and videos in the PDID are authentic or fake. We code 1) whether the sharer presented the image or video as real or fake, and 2) any external verification of veracity. Coders also identify visual markings that could indicate content is synthetic (e.g., extra fingers, blurred faces).

Context and Content: We collect information regarding the context and content of the deepfakes, using text or transcripts from the image or video itself as well as from the primary associated post and surrounding comments. We further identify *communication goals* for each deepfake, encompassing a range of objectives from political interference, emotional and psychological harm, and reputational damage to satire, entertainment, and education. That is, while deepfakes are often associated with spreading false information, our coding strategy acknowledges their potential for goals like education or raising awareness (Schegloff 1997; Poria et al. 2017). Additionally, we consider potential harms that political deepfakes might *depict*. Unlike existing frameworks focusing on physical harm, we develop a harm type typology that encompasses a broader range of harms. In particular, we identify not only whether political deepfakes depict harms associated with violence or property damage, but

also whether they depict human rights violations, discrimination, or even environmental harms. Communication goals and harm types are not mutually exclusive, and we aim for a comprehensive understanding of deepfakes’ multifaceted dangers. Determining communication goals and harm types involves subjectivity and inference, but is crucial for understanding these complex social phenomena.

Real-world Connections: Next, the PDID aims to connect political deepfakes to *real-world events and harms* to better understand the motivation for their creation and spread. We identify whether each deepfake is associated with contemporaneous events such as scandals, crimes, elections, or wars. We also code whether there is potential evidence, based on news reports, of actual harms tied to the deepfakes. For example, as described previously, a fake image of a bombing at the Pentagon caused financial harm through a dip in the stock market. We thus identify whether political deepfakes have caused harms related to financial loss, political interference, discrimination, human rights, violence, property damage, or the environment. For transparency, coders provide information regarding the evidence for their inferences, such as links to external reporting.

Politics and Policy: Given that the PDID includes political deepfakes, we also devote particular attention to factors related to politics and policy. We categorize the *policy sectors* most relevant to each deepfake using definitions from the CSET policy sector typology and the U.S. Census Bureau (US Census Bureau 2022). We distinguish between sectors like medicine, transportation, and public safety/military, and we allow deepfakes to fall into multiple policy sectors. This helps to identify intended targets and potential impacts, revealing patterns in sectors most affected by or vulnerable to deepfakes. Next, we extend *framing* research, which typically focuses on the media (Matthes 2009) and framing efforts by actors like advocacy groups (De Bruycker 2017), to deepfakes and the actors sharing them. The framing of deepfakes refers to the perspectives, themes, or underlying ideas that shape how the manipulated content is communicated. We draw on Semetko and Valkenburg (2000)’s typology of human interest, conflict, economic consequences, morality, and responsibility frames, which helps to uncover underlying motivations, such as intentions to exacerbate conflicts or assign blame. Finally, we identify aspects of the *policy narrative* represented within each deepfake. This includes who or what the deepfake portrays as the hero, villain, plot, and moral, building on research on the Narrative Policy Framework (NPF) in the policy process literature (Shanahan, McBeth, and Hathaway 2011). The study of narratives allows for investigations into how deepfakes impact individuals’ understanding of policy actors, underlying responsibility, and social problems and solutions.

Conclusion

Opportunities for Research and Practice: The PDID provides numerous opportunities to address harms and advance governance efforts related to political deepfakes. From a research perspective, the database is well-suited for longitudinal studies, surveys, experiments, and descriptive analyses.

⁷Available online at DOI 10.17605/OSF.IO/FVQG3.

Analysts focused on political actors, institutions, and the political economy of misinformation can examine which actors are employing deepfakes and for what purposes. Further, the PDID can be used to analyze the prevalence of different frames or narratives used in political deepfakes, shedding light on the strategies employed to shape public perceptions. The database can thus facilitate research on how exposure to political deepfakes influences public opinion, trust in media, and other key aspects of political behavior, such as voting.

The database can also advance research on ethical considerations surrounding AI, the creation and distribution of political deepfakes, and legal implications related to privacy, defamation, image-based sexual violence, and propaganda. Researchers and practitioners can use the database to investigate how political deepfakes impact public discourse, democratic processes, and civic engagement. Importantly, within each of the domains described above, researchers could also conduct comparative studies across different political contexts and demographic subgroups to explore how the use of and responses to political deepfakes vary.

The PDID can also support researchers, practitioners, and policymakers working on AI policy and technology governance. For example, it can enable exploration into how policy changes or advancements in deepfake technologies influence the distribution and content of political deepfakes over time. Did the public release of large language models like ChatGPT change how AI-generated political content is disseminated? Can the addition of watermarks to deepfakes successfully increase detection or decrease their spread? Can enhanced transparency requirements, incident reporting, fact-checking, or auditing measures counteract harms from deepfake proliferation or improve public resilience?

Importantly then, the PDID can aid in developing and evaluating countermeasures, fact-checking efforts, detection techniques, and policy actions by AI developers, social media companies, standards organizations, and regulators. This includes informing the development of media or digital literacy initiatives to equip individuals with the skills to identify and critically assess political deepfakes. A central purpose of the PDID, in linking AI incidents research with misinformation research, is to help actualize an applied research agenda to address these pressing societal challenges.

Future Directions: Nonetheless, realizing these goals will require significant effort to enhance, govern, and sustain the PDID. Key goals and associated challenges include:

1. **Maintaining and sustaining the database.** The PDID is intended to adopt FAIR principles (Findable, Accessible, Interoperable, and Reusable), including preserving an associated downloaded file for each deepfake. Pending consideration of increased ethical risks, the PDID will be an open-access resource, easily available for researchers, policymakers, journalists, and citizens. This requires sustainable infrastructure, funding, and staffing.
2. **Enabling public incident reporting.** This will facilitate scaling, shared governance, and attention to incidents that could otherwise be neglected, and requires infrastructures for both reporting and vetting reports.
3. **Promoting ecosystem interoperability.** We aim to

make the PDID interoperable with other AI incident and vulnerability reporting tools. Ongoing conversations with AIID and AVID, and future engagement with government-sanctioned incident reporting, will improve the impact and reach of the PDID and related efforts.

4. **Expanding regional and linguistic coverage.** Language and regional coverage limitations represent a significant gap in the current PDID, exacerbating a dilemma that individuals in low-income regions can suffer from both more misinformation and less mitigation. The PDID will benefit from efforts to extend data collection to other platforms, communication methods, countries, and languages. This will require careful scaling, local expertise, political sensitivity, and multilayered governance.
5. **Developing shared governance.** The PDID will only be effective if it is actually used by researchers and other key practitioners. We will thus continue efforts to engage with journalists, fact-checking organizations, policymakers, researchers, and civil society groups. Governance mechanisms could include, for example, a steering committee and ongoing forums for feedback to enhance the utility of the PDID for different stakeholders.
6. **Expanding dataset coverage.** As policy developments and regulations impact how deepfakes are produced and disseminated (e.g., through the inclusion of watermarks or formal verification efforts), we will iterate on the breadth and depth of the coding categories in the PDID to provide a more robust, fine-grained, and comprehensive database. We will also explore automated techniques to improve coverage and efficiency compared to manual coding. Finally, expanding the PDID to encompass other media formats like audio may be socially important.

Ethical Statement. Curating and storing deepfakes has the potential to increase access to harmful content. However, we believe that the PDID’s record of historical and verified (when possible) deepfakes is unlikely to be useful for adversarial purposes above and beyond existing tools, while offering new benefits for research, verification, and mitigation. Given concerns about psychological harms from exposure to deepfakes, the PDID excludes highly incendiary sexual or violent content. We will continue to weigh the benefits and risks of open access data on political deepfakes in decisions about the public release and presentation of the PDID.

Acknowledgments

The authors would like to thank the Center for Security and Emerging Technology (CSET) for access to their harms taxonomy and codebook on AI Incidents, which supported the development of the PDID codebook. We are grateful to the AI Incidents Database (AIID) and AI Vulnerabilities Database (AVID) teams for their insights and feedback on interoperability. We also would like to thank J.P. Messina for helping to develop the codebook and broader project agenda, and are grateful to our undergraduate research assistants Paige Carter, Sharvani Kondapally, and Shreya Venkat.

References

- Alba, D. 2023. Midjourney Is Easily Tricked Into Making AI Misinformation, Study Finds. *Bloomberg*.
- Barari, S.; Lucas, C.; and Munger, K. 2021. Political Deepfakes Are As Credible As Other Fake Media And (Sometimes) Real Media. Preprint, Open Science Framework.
- Cahlan, S. 2020. Analysis | How misinformation helped spark an attempted coup in Gabon. *Washington Post*.
- Christopher, N. 2020. We've Just Seen the First Use of Deepfakes in an Indian Election Campaign.
- Cook, E. 2019. Deep fakes could have real consequences for Southeast Asia.
- Dao, D.; W, H.; Walla, A.-A.; VocalFan; Rubiel, E.; Schreiber, F.; Jaworski, J.; Liu, L.; Williams, N.; Pawlowski, N.; Ammanamanchi, P. S.; Stadlmann, S.; Kühne, S.; Taiz; Diethe, T.; jvmncs; and twsl. 2022. Awful AI.
- De Bruycker, I. 2017. Framing and advocacy: A research agenda for interest group studies. *Journal of European Public Policy*, 24(5): 775–787.
- European Commission. 2021. Proposal for a regulation on a European approach for artificial intelligence. Technical report, European Commission, Brussels, Belgium.
- Goldstein, J. A.; Sastry, G.; Musser, M.; DiResta, R.; Gentzel, M.; and Sedova, K. 2023. Forecasting Potential Misuses of Language Models for Disinformation Campaigns—and How to Reduce Risk. Technical report, Brookings Institution.
- Groh, M.; Epstein, Z.; Firestone, C.; and Picard, R. 2022. Deepfake Detection by Human Crowds, Machines, and Machine-informed Crowds. *Proceedings of the National Academy of Sciences*, 119(1): 1–11.
- Hao, K. 2023. China, a Pioneer in Regulating Algorithms, Turns Its Focus to Deepfakes. *Wall Street Journal*.
- Klar, R. 2023. Democrats laud federal agency's AI move. *The Hill*.
- Matthes, J. 2009. What's in a Frame? A Content Analysis of Media Framing Studies in the World's Leading Communication Journals, 1990-2005. *Journalism & Mass Communication Quarterly*, 86(2): 349–367.
- McGregor, S. 2021. Preventing Repeated Real World AI Failures by Cataloging Incidents: The AI Incident Database. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17): 15458–15463.
- McGregor, S.; Paeth, K.; and Lam, K. 2022. Indexing AI Risks with Incidents, Issues, and Variants. *arXiv preprint*.
- Montanaro, D. 2022. The truth about political ads: They can include lies. *NPR*.
- Nehamas, N. 2023. DeSantis Campaign Uses Apparently Fake Images to Attack Trump on Twitter. *The New York Times*.
- Partnership on AI. 2023. PAI's Responsible Practices for Synthetic Media. Technical report, Partnership on AI.
- Poria, S.; Cambria, E.; Hazarika, D.; Majumder, N.; Zadeh, A.; and Morency, L.-P. 2017. Context-Dependent Sentiment Analysis in User-Generated Videos. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 873–883. Vancouver, Canada: Association for Computational Linguistics.
- Rana, M. S.; Nobi, M. N.; Murali, B.; and Sung, A. H. 2022. Deepfake Detection: A Systematic Literature Review. *IEEE access : practical innovations, open solutions*, 10: 25494–25513.
- Schegloff, E. A. 1997. Whose Text? Whose Context? *Discourse & Society*, 8(2): 165–187.
- Schiff, K. J.; Schiff, D. S.; and Bueno, N. 2023. The Liar's Dividend: Can Politicians Claim Misinformation to Evade Accountability? *SocArXiv*.
- Semetko, H. A.; and Valkenburg, P. M. V. 2000. Framing European politics: A Content Analysis of Press and Television News. *Journal of Communication*, 50(2): 93–109.
- Shanahan, E. A.; McBeth, M. K.; and Hathaway, P. L. 2011. Narrative policy framework: The influence of media policy narratives on public opinion. *Politics & Policy*, 39(3): 373–400.
- Sharma, I.; Jain, K.; Behl, A.; Baabdullah, A.; Giannakis, M.; and Dwivedi, Y. 2023. Examining the motivations of sharing political deepfake videos: The role of political brand hate and moral consciousness. *Internet Research*, ahead-of-print(ahead-of-print).
- Simonite, T. 2022. A Zelensky Deepfake Was Quickly Defeated. The Next One Might Not Be. *Wired*.
- Sorkin, A. R.; Warner, B.; Kessler, S.; Merced, M. J. d. I.; Hirsch, L.; and Livni, E. 2023. An A.I.-Generated Spoof Rattles the Markets. *The New York Times*.
- Stokols, E.; and Egan, L. 2023. White House press shop adjusts to proliferation of AI deep fakes. *POLITICO*.
- US Census Bureau. 2022. Economic Census: NAICS Codes & Understanding Industry Classification Systems. Technical report.
- Vaccari, C.; and 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): 2056305120903408.
- Williams, K. 2023. Exploring Legal Approaches to Regulating Nonconsensual Deepfake Pornography. *Tech Policy Press*.
- Wittenberg, C.; Tappin, B. M.; Berinsky, A. J.; and Rand, D. G. 2021. The (minimal) persuasive advantage of political video over text. *Proceedings of the National Academy of Sciences*, 118(47).