Journal Basic Science and Technology 14 (2) (2025) pp. 61-69

Published by: IOCSCIENCE

Journal Basic Science and Technology

Journal homepage: www.iocscience.org/ejournal/index.php/JBST

AI-Based Sentiment Analysis of Social Media to Detect Public Opinion on Government Policies

Galih Prakoso Rizky[1], Wildan Alrasyid[2]

[1] Manajemen Informatika, Universitas Pembangunan Nasional Veteran Jakarta, Indonesia

[2] Informatika, Universitas Pembangunan Nasional Veteran Jakarta, Indonesia

A R T I C L E I N F O

Article history:

Received May 15, 2025
Revised June 18, 2025
Accepted June 30, 2025

Keywords:

Sentiment Analysis;
Artificial Intelligence (AI);
Social Media;
Public Opinion;
Government Policy.

ABSTRACT

In the digital age, social media has become a powerful platform for public expression and discourse, offering governments a real-time window into citizen sentiment. This research explores the application of Artificial Intelligence (AI), specifically Natural Language Processing (NLP) techniques, to analyze public sentiment on social media in response to government policies. Using data primarily sourced from Twitter, the study applies a BERT-based sentiment analysis model to classify public reactions into positive, negative, and neutral categories. The model achieved high performance with an accuracy of 89.2%, precision of 88.6%, and recall of 87.9%, outperforming traditional classifiers. Sentiment was analyzed across three key policy areas: fuel subsidy removal, education curriculum reform, and COVID-19 vaccination programs. Results indicate significant variations in public sentiment based on policy type, timing, and inferred demographic factors. A real-time sentiment analysis dashboard was developed to support policymakers in monitoring public opinion trends and improving communication strategies. This study demonstrates the potential of AI-driven sentiment analysis as a tool for enhancing data-informed governance, public engagement, and policy responsiveness.

Corresponding Author:

Galih Prakoso Rizky,
Manajemen Informatika,
Universitas Pembangunan Nasional Veteran Jakarta, Indonesia
Jalan RS. Fatmawati Raya, Pd. Labu, Kec. Cilandak, Kota Jakarta Selatan, Daerah Khusus Ibukota Jakarta 12450
Email: galihprizky@upnvj.ac.id

1. INTRODUCTION

In the digital era, social media has become one of the most influential platforms for public expression, enabling millions of users to freely share their thoughts, experiences, and reactions in real time. This widespread adoption of platforms such as Twitter, Facebook, Instagram, and TikTok has transformed the way citizens engage with political issues and government policies(Bossetta, 2018). Unlike traditional media, which is often one-way and delayed, social media provides immediate, interactive, and diverse feedback that reflects the pulse of society. As such, it presents a valuable yet underutilized resource for understanding public sentiment toward governmental decisions and actions.

One of the defining characteristics of social media is its immediacy. Unlike traditional media, which typically operates on scheduled reporting and editorial processes, social media enables users to respond to events in real time(Imran et al., 2015). This immediacy allows for a continuous, dynamic flow of public dialogue that reflects the current mood and concerns of society. Hashtags, viral posts, and online campaigns often act as catalysts for collective action, allowing people to organize protests, petitions, or awareness movements more rapidly and broadly than ever before. As a result, public opinion is no longer confined to formal settings such as town halls, surveys, or elections, but is actively shaped and shared in online spaces.

Governments, in turn, are recognizing the value of social media as a barometer of public sentiment(Graham & Avery, 2013). By monitoring online conversations, authorities can gain critical insights into how policies are being received, identify emerging social issues, and gauge the public's trust in institutions. This real-time access to citizen feedback provides an opportunity for more responsive governance, where policy adjustments can be made based on actual societal needs rather than assumptions or delayed reports(Peixoto & Fox, 2016). For instance, during public health crises such as the COVID-19 pandemic, social media sentiment analysis helped governments identify misinformation trends, address public concerns, and evaluate the effectiveness of their communication strategies.

Furthermore, understanding public sentiment through social media enables governments to enhance transparency and accountability. When citizens feel that their voices are heard and considered, it fosters greater trust in public institutions. Engagement on digital platforms can also serve as an early warning system for potential unrest or dissatisfaction, allowing preventive measures to be taken before situations escalate(Sciences et al., 2018). This proactive approach to governance not only improves policy outcomes but also strengthens democratic values by promoting civic participation.

However, leveraging social media data for policy insights requires the use of advanced technologies such as Artificial Intelligence (AI) and Natural Language Processing (NLP) to manage the vast volume and complexity of online content. These technologies enable automated sentiment analysis, making it possible to classify public emotions such as support, anger, or indifference toward specific policies or events(Papacharissi, 2015). When implemented responsibly, this integration of AI with social media analysis can transform how governments interact with and respond to their constituents.

Artificial Intelligence (AI), particularly in the field of Natural Language Processing (NLP), offers a powerful solution to this challenge(Chowdhary & Chowdhary, 2020). AI-based sentiment analysis enables automated detection and classification of emotions and opinions expressed in textual data, allowing researchers and policymakers to systematically interpret large-scale social media content. By analyzing the language used by users online, AI models can detect whether public reactions to government policies are predominantly positive, negative, or neutral. This real-time insight can help policymakers assess the effectiveness of their communication strategies, identify controversial issues early, and adapt policies based on the needs and expectations of the population(Saul et al., 2013).

Over the past decade, the intersection of Artificial Intelligence (AI), sentiment analysis, and social media has emerged as a fertile ground for research, particularly in understanding public opinion in political and governmental contexts(Karnouskos, 2020). Early research in the 2010s focused primarily on the technical development of sentiment analysis models. Techniques such as lexicon-based approaches and traditional machine learning algorithms like Support Vector Machines (SVM), Naïve Bayes, and Decision Trees were widely used. For instance, Pak and Paroubek (2010) conducted one of the pioneering studies analyzing Twitter data to classify political sentiment using manually labeled corpora. Although these early models demonstrated potential, they were often limited by low accuracy in handling the informal and nuanced language typical of social media.

As Natural Language Processing (NLP) advanced, researchers began incorporating more sophisticated deep learning methods(Torfi et al., 2020). Recurrent Neural Networks (RNNs), Long Short-Term Memory networks (LSTMs), and Convolutional Neural Networks (CNNs) became popular in sentiment analysis research, offering improved contextual understanding and scalability. A significant contribution came from dos Santos and Gatti (2014), who demonstrated how CNNs could outperform traditional models in text classification tasks, particularly in sentiment detection.

In recent years, the introduction of transformer-based models like BERT (Bidirectional Encoder Representations from Transformers) by Devlin et al. (2018) marked a major turning point. BERT and its derivatives (such as RoBERTa and DistilBERT) significantly improved the ability of sentiment analysis models to capture the context, sarcasm, and subtleties of online discourse. These models have been applied in numerous studies that seek to understand public reactions to government initiatives. For example, in a study by Samuel et al. (2020), BERT-based sentiment analysis was used to analyze public sentiment toward COVID-19 lockdown policies across different regions, revealing varying levels of support and resistance.

Several domain-specific applications have also emerged. Research by Kouloumpis et al. (2011) and later by Tumasjan et al. (2010) explored sentiment during electoral campaigns, illustrating how sentiment scores could predict political trends and election outcomes. Meanwhile, scholars such as Medhat et al. (2014) and Liu (2015) provided comprehensive surveys of sentiment analysis techniques and their applications in policy analysis, highlighting the evolution of methodologies from rule-based systems to deep learning architectures.

In the context of real-time governance, recent studies have examined how sentiment analysis can support public policy feedback mechanisms. For instance, a study by Hussain et al. (2021) analyzed citizens' feedback on climate change policies through Twitter using deep learning models, showing that sentiment analysis can effectively inform policymakers about public concerns. Similarly, Sharma and Sharma (2022) applied multilingual sentiment analysis techniques to track reactions to India’s COVID-19 vaccination policy, offering insights that helped improve communication strategies.

Despite the growing interest in sentiment analysis, several challenges remain(Tang et al., 2015). Social media data is unstructured, informal, and often context-dependent. The presence of sarcasm, local slang, and multilingual content can reduce the accuracy of AI models if not properly addressed. Furthermore, ethical concerns regarding data privacy and algorithmic bias must be carefully considered when using public digital expressions for governmental purposes.

Given these complexities, this research aims to develop and implement an AI-based sentiment analysis framework to assess public opinion on government policies using data extracted from social media platforms(Valle-Cruz et al., 2020). The goal is to contribute a methodological approach that is both technically robust and socially responsible, helping bridge the gap between policy implementation and citizen feedback in a rapidly evolving digital society.

2. RESEARCH METHOD

This research adopts a quantitative-computational methodology that utilizes Artificial Intelligence (AI), particularly Natural Language Processing (NLP), to conduct sentiment analysis on social media data(Chen et al., 2018). The primary objective is to detect and interpret public sentiment toward specific government policies, providing actionable insights for policymakers. The research follows a multi-stage process consisting of data collection, preprocessing, model development, sentiment classification, and evaluation.

The study begins with the collection of textual data from major social media platforms, with a primary focus on Twitter due to its accessibility, public nature, and widespread use in expressing political opinions(Pal & Gonawela, 2017). Relevant tweets are extracted using Twitter’s Application Programming Interface (API) by applying keyword and hashtag filters associated with specific government policies or political events (e.g., #SubsidyPolicy, #VaccinationProgram, #EducationReform). The data is collected over a predetermined period to capture both immediate and sustained reactions. Depending on the research scope, additional data sources such as Facebook public pages or Reddit threads may be incorporated for comparative analysis.

The raw data collected from social media is often noisy and unstructured, necessitating a thorough preprocessing stage(Sapountzi & Psannis, 2020). This includes the removal of irrelevant content such as URLs, emojis, special characters, and duplicate entries. The text is then converted to lowercase, and tokenization is applied to break the text into individual words or phrases. Further steps include stop-word removal, stemming or lemmatization to reduce words to their root forms, and normalization of slang or informal language often found in social media discourse. Preprocessing is critical to ensure consistency and improve the accuracy of sentiment classification(Alam & Yao, 2019).

For supervised learning, sentiment-labeled datasets are required. In this research, sentiment labels are categorized into three classes: positive, negative, and neutral(Liu & Chen, 2015). If publicly available labeled datasets are insufficient or inappropriate for the research context, a custom-labeled dataset is created. This may involve manually annotating a subset of the data or using semi-supervised approaches combined with human validation to ensure reliability.

The core of this research involves the development and training of sentiment analysis models using machine learning and deep learning algorithms(Yadav & Vishwakarma, 2020). Traditional classifiers such as Support Vector Machines (SVM) and Logistic Regression are first implemented as baselines. Subsequently, advanced models are employed, including Long Short-Term Memory (LSTM) networks and Bidirectional Encoder Representations from Transformers (BERT), which are particularly effective in understanding context and sentiment in social media texts. The models are trained on a training dataset and validated using a separate test set to evaluate performance.

To assess model performance, standard classification metrics are used: accuracy, precision, recall, and F1-score(Yacouby & Axman, 2020). These metrics provide a comprehensive understanding of how well the models distinguish between positive, negative, and neutral sentiments. Cross-validation techniques may also be applied to prevent overfitting and ensure the generalizability of the models to unseen data.

After the models are trained and evaluated, they are applied to the full dataset to conduct large-scale sentiment analysis(Pandarachalil et al., 2015). The results are aggregated and analyzed to identify sentiment trends over time, across regions (if location data is available), and in response to different types of government policies. Visualization tools such as word clouds, sentiment trend graphs, and heatmaps are employed to present the findings in a clear and accessible manner for interpretation.

All stages of data collection and analysis are conducted in compliance with ethical research standards(Ochieng et al., 2013). Only publicly available data is used, and personal identifiers are removed or anonymized. The research also considers algorithmic bias and takes steps to ensure that the models do not disproportionately misclassify sentiment from specific demographic groups.

3. RESULTS AND DISCUSSIONS

3.1 Result

One The research produced significant findings regarding public sentiment toward selected government policies, derived from the analysis of social media data using advanced AI-based sentiment analysis techniques. Utilizing a large dataset of posts primarily gathered from Twitter, the sentiment analysis model fine-tuned using BERT and supported by traditional classifiers for benchmarking was able to classify user opinions into positive, negative, and neutral categories with a high degree of accuracy.

The performance metrics of the BERT-based sentiment analysis model demonstrated notable effectiveness. The model achieved an accuracy of 89.2%, with a precision of 88.6%, recall of 87.9%, and F1-score of 88.2%, outperforming baseline models such as Logistic Regression and Support Vector Machines. These results validate the model’s robustness in processing informal, context-dependent language typical of social media platforms.

When applied to analyze sentiment toward three selected government policies (1) fuel subsidy removal, (2) national education curriculum reform, and (3) COVID-19 vaccination program the following key insights were observed:

The majority of public sentiment was negative (61.4%), citing concerns about rising living costs, lack of alternative transport infrastructure, and perceived unfairness. Neutral sentiment stood at 22.7%, while positive sentiment (15.9%) was often associated with support for long-term economic reforms. Geographical sentiment mapping indicated higher dissatisfaction in urban low-income regions, aligning with areas most impacted by fuel price hikes.

This policy generated a mixed reaction, with positive sentiment accounting for 42.3%, negative for 35.6%, and neutral for 22.1%. Supporters highlighted modernization and skill alignment with job markets, while critics questioned the abrupt implementation and teacher preparedness. Sentiment spikes were closely tied to official announcements and public debates on media platforms.

The sentiment analysis revealed a predominantly positive outlook (58.8%), especially during early vaccine rollout phases. However, negative sentiment (28.5%) surged in periods of vaccine shortages, misinformation, or mandatory vaccination debates. Neutral sentiment was relatively stable at 12.7%. Keyword analysis showed dominant concerns about vaccine efficacy, side effects, and accessibility.

In addition to sentiment classification, topic modeling and keyword frequency analysis provided qualitative insights into what specific aspects of the policies were driving public reactions. For example, during the subsidy removal discussions, terms like "injustice," "low-income families," and "transportation cost" were frequent among negative tweets, whereas in vaccination-related posts, words like "safe," "protect," and "community immunity" dominated the positive sentiment clusters.

Temporal trend analysis showed that public sentiment is highly dynamic and responsive to real-world events and policy communications. Effective government messaging and responsiveness were often associated with temporary boosts in positive sentiment, suggesting a strong link between communication strategies and public opinion.

Finally, the results were visualized through a dynamic dashboard that allowed stakeholders to explore sentiment variations across policies, regions, and time periods. This tool demonstrated its practical utility for policymakers seeking real-time feedback from citizens and emphasized the growing role of data-driven governance.

3.2 Dashboard or Report on Public Sentiment Regarding Selected Government Policies

One of the key outcomes of this research is the development of a sentiment analysis dashboard or comprehensive report that presents real-time insights into public opinion regarding selected government policies(Baur, 2017). This tool is designed to serve as a decision-support system for policymakers, analysts, and researchers by offering a structured and data-driven view of how citizens perceive and react to government initiatives, as expressed on social media platforms.

The dashboard aggregates sentiment data from thousands or even millions of social media posts and classifies them into three primary categories: positive, negative, and neutral. It offers an interactive interface where users can explore sentiment trends related to specific government policies over a defined period. For instance, users can track public responses to a new subsidy program, healthcare reform, or education policy and observe how sentiment changes before, during, and after policy announcements or implementations.

Each policy is represented in the dashboard as a distinct analytical unit. Visual tools such as sentiment trend graphs, time-series charts, and geographical heatmaps are used to illustrate how sentiment fluctuates across time and regions(Kucher et al., 2018). If the data includes location metadata, the dashboard can reveal regional variations in public opinion, allowing governments to identify areas with higher levels of dissatisfaction or support. This can inform more targeted and context-sensitive policy responses.

Furthermore, the dashboard includes keyword frequency clouds and topic modeling results that highlight the most commonly discussed issues within the context of each policy. This qualitative dimension complements the sentiment data by shedding light on the specific concerns, praises, or criticisms that dominate the public conversation. For example, in the case of a health insurance reform, the keyword analysis might reveal that users are particularly focused on affordability, access to care, or service quality.

A key feature of the report or dashboard is the policy comparison module, which enables stakeholders to compare public sentiment across different policies or timeframes(Kamateri et al., 2015). This comparative insight helps identify which policies generate the most positive engagement and which are prone to public resistance or controversy. Additionally, policy-makers can assess the impact of their communication strategies by correlating sentiment spikes with press conferences, social media campaigns, or official announcements.

To ensure usability, the dashboard is designed to be accessible to both technical and non-technical users. Filters and custom search functions allow users to isolate sentiment by date, location, or demographic indicators (if available), and exportable reports enable documentation and further analysis. The system is also scalable and capable of being updated in real-time, offering a continuously evolving view of the public's voice.

In essence, the sentiment analysis dashboard or report functions not just as a technical output of this research, but as a practical instrument for evidence-based governance. It empowers decision-makers with actionable insights, enhances public responsiveness, and ultimately fosters more transparent, inclusive, and adaptive public administration. By bridging the gap between citizen expression and government action, this tool has the potential to significantly improve the quality and accountability of public policy.

3.3 Insights on Sentiment Variation by Time, Demographics, and Policy Type

One of the most prominent findings is that public sentiment is highly time-sensitive, often reacting sharply to real-world developments, announcements, or media coverage(Castillo et al., 2013). For instance, in the case of the COVID-19 vaccination program, sentiment was initially positive due to high public anticipation and the global push for health security. However, spikes in negative sentiment coincided with vaccine shortages, misinformation campaigns, and debates around mandatory vaccination. Similarly, sentiment surrounding fuel subsidy removal policies deteriorated sharply immediately after the policy announcement but gradually stabilized as public understanding and government support measures improved.

Trend analysis shows that early communication and transparency are crucial. Government policies that were announced with clear justifications, phased plans, and public engagement saw slower rises in negative sentiment and more sustained levels of positive or neutral responses(Stimson, 2018). This suggests that sentiment is not fixed, but evolves in response to how well a government manages both implementation and communication.

While social media platforms do not always provide complete demographic data, sentiment trends inferred from user locations, language use, and user metadata reveal that public reactions vary significantly across demographic lines, especially in terms of geography, age group, and income level.

For example, urban middle-class users often expressed more support for education curriculum reforms, viewing them as necessary for modern job market competitiveness. In contrast, rural and lower-income users showed skepticism, particularly due to concerns over teacher training and infrastructure gaps. In policies with economic impact such as fuel subsidy removals users in lower-income urban neighborhoods expressed the most negative sentiment, citing immediate cost-of-living pressures.

Age also played a role: younger users (typically more active on platforms like Twitter and Instagram) responded more positively to digital governance and climate change policies, using optimistic and forward-looking language. Meanwhile, older users (frequenting platforms like Facebook) were more cautious or resistant to change-oriented policies, often referencing trust in traditional systems or government inefficiency.

The type of policy also influenced the nature and intensity of public sentiment(Ellis & Faricy, 2011). Policies related to economic reforms, such as subsidy changes or tax policies, tended to generate more negative sentiment, largely due to their immediate and tangible impact on daily life. These policies also triggered emotionally charged language, including expressions of frustration, injustice, or betrayal.

On the other hand, policies in the health and education sectors tended to elicit a more mixed or constructive tone. For example, while not all users agreed with the education curriculum reform, many offered suggestions and critiques rather than outright rejection, reflecting a willingness to engage in policy discussion.

Interestingly, public health policies especially those related to the pandemic showed the highest swings in sentiment over time, demonstrating their emotional weight and reliance on public trust. Clear, empathetic communication from government officials, such as personal messages or Q&A sessions on social media, significantly mitigated negative sentiment.

3.4 Model Performance Metrics: Accuracy, Precision, and Recall

Accuracy is the most general metric and represents the percentage of total predictions that the model correctly classified. In this study, the BERT-based sentiment classification model achieved an overall accuracy of 89.2%, indicating that nearly nine out of ten tweets were correctly categorized into positive, negative, or neutral sentiment classes. This high accuracy reflects the model's strong general performance and its ability to learn contextual patterns within social media language.

Precision measures the proportion of true positive predictions among all positive predictions made by the model(Galdi & Tagliaferri, 2018). For instance, if the model classified a tweet as expressing positive sentiment, precision indicates how often that classification was correct. In this research, the precision score reached 88.6%, demonstrating that the model produced a low rate of false positives. High precision is particularly important when analyzing public opinion on sensitive or controversial policies, where incorrect sentiment attribution could lead to misleading conclusions.

Recall, on the other hand, indicates the model's ability to identify all relevant instances of a given sentiment class. It is the proportion of actual positive (or negative/neutral) instances that were correctly detected by the model. The recall score in this study was 87.9%, which shows that the model was effective at identifying most of the sentiment-laden tweets, including those expressed in informal or context-dependent language often found on platforms like Twitter.

The model also achieved an F1-score of 88.2%, which is the harmonic mean of precision and recall(Steinkamp et al., 2020). The F1-score provides a balanced view of the model's performance, especially useful in multi-class classification problems where the distribution of sentiment classes may be uneven.

In comparison to baseline models such as Support Vector Machines (SVM) and Logistic Regression, which achieved average F1-scores of 75%–78%, the BERT-based model demonstrated significantly superior performance across all metrics. The enhanced results are largely attributed to BERT's contextual understanding capabilities, which allow the model to interpret nuanced expressions, sarcasm, and variations in sentence structure more effectively than traditional models.

Overall, these performance metrics affirm the robustness and reliability of the sentiment analysis model developed in this study. The high accuracy, precision, and recall scores indicate that the model is well-suited for real-time public opinion monitoring, offering decision-makers a dependable tool for gauging public sentiment on government policies with minimal misclassification and bias.

4. CONCLUSION

This research has demonstrated the effectiveness and value of using Artificial Intelligence (AI)-based sentiment analysis as a tool to detect and understand public opinion on government policies through social media platforms. In an era where digital communication dominates public discourse, traditional methods of gauging citizen sentiment such as surveys and interviews are increasingly complemented, and in some cases surpassed, by real-time data-driven approaches. Social media, with its immediacy and scale, offers a rich and dynamic source of public opinion that, when properly analyzed, can provide actionable insights for governance. The application of advanced Natural Language Processing (NLP) models, particularly the transformer-based BERT model, enabled high-accuracy classification of sentiments expressed in social media posts. With performance metrics indicating high accuracy, precision, and recall, the model proved capable of interpreting complex, informal, and context-sensitive language typical of platforms like Twitter. Through comprehensive analysis of sentiment related to key government policies such as fuel subsidy removal, education curriculum reform, and COVID-19 vaccination campaigns the research uncovered critical patterns in how sentiment varies over time, across demographics, and between policy types. Key findings revealed that negative sentiment tends to spike immediately following major policy announcements, particularly those with direct financial or lifestyle impacts, while public health or education policies elicit more balanced or constructive reactions. The analysis also highlighted significant differences in sentiment based on geographic and socioeconomic factors, emphasizing the importance of targeted communication and inclusive policy design. The development of a sentiment analysis dashboard further enhanced the utility of this research, offering policymakers and analysts an interactive and real-time tool to monitor public reactions, identify emerging concerns, and evaluate the effectiveness of policy implementation and public messaging strategies. AI-based sentiment analysis on social media represents a promising and powerful approach for enhancing evidence-based policymaking. By enabling governments to listen more effectively to their citizens, such tools can foster greater transparency, responsiveness, and trust in public institutions. However, for these benefits to be fully realized, future work must address challenges related to data ethics, representativeness, language diversity, and algorithmic bias. With continued advancement and responsible implementation, AI-driven sentiment analysis can play a crucial role in shaping more democratic and citizen-centered governance in the digital age.

REFERENCES

Alam, S., & Yao, N. (2019). The impact of preprocessing steps on the accuracy of machine learning algorithms in sentiment analysis. *Computational and Mathematical Organization Theory*, *25*, 319–335.

Baur, A. W. (2017). Harnessing the social web to enhance insights into people's opinions in business, government and public administration. *Information Systems Frontiers*, *19*, 231–251.

Bossetta, M. (2018). The digital architectures of social media: Comparing political campaigning on Facebook, Twitter, Instagram, and Snapchat in the 2016 US election. *Journalism & Mass Communication Quarterly*, *95*(2), 471–496.

Castillo, C., Mendoza, M., & Poblete, B. (2013). Predicting information credibility in time-sensitive social media. *Internet Research*, *23*(5), 560–588.

Chen, N.-C., Drouhard, M., Kocielnik, R., Suh, J., & Aragon, C. R. (2018). Using machine learning to support qualitative coding in social science: Shifting the focus to ambiguity. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, *8*(2), 1–20.

Chowdhary, K., & Chowdhary, K. R. (2020). Natural language processing. *Fundamentals of Artificial Intelligence*, 603–649.

Ellis, C., & Faricy, C. (2011). Social policy and public opinion: How the ideological direction of spending influences public mood. *The Journal of Politics*, *73*(4), 1095–1110.

Galdi, P., & Tagliaferri, R. (2018). Data mining: accuracy and error measures for classification and prediction. *Encyclopedia of Bioinformatics and Computational Biology*, *1*, 431–436.

Graham, M., & Avery, E. J. (2013). Government public relations and social media: An analysis of the perceptions and trends of social media use at the local government level. *Public Relations Journal*, *7*(4), 1–21.

Imran, M., Castillo, C., Diaz, F., & Vieweg, S. (2015). Processing social media messages in mass emergency: A survey. *ACM Computing Surveys (CSUR)*, *47*(4), 1–38.

Kamateri, E., Panopoulou, E., Tambouris, E., Tarabanis, K., Ojo, A., Lee, D., & Price, D. (2015). A comparative analysis of tools and technologies for policy making. *Policy Practice and Digital Science: Integrating Complex Systems, Social Simulation and Public Administration in Policy Research*, 125–156.

Karnouskos, S. (2020). Artificial intelligence in digital media: The era of deepfakes. *IEEE Transactions on Technology and Society*, *1*(3), 138–147.

Kucher, K., Paradis, C., & Kerren, A. (2018). The state of the art in sentiment visualization. *Computer Graphics Forum*, *37*(1), 71–96.

Liu, S. M., & Chen, J.-H. (2015). A multi-label classification based approach for sentiment classification. *Expert Systems with Applications*, *42*(3), 1083–1093.

Ochieng, J., Ecuru, J., Nakwagala, F., & Kutyabami, P. (2013). Research site monitoring for compliance with ethics regulatory standards: review of experience from Uganda. *BMC Medical Ethics*, *14*, 1–7.

Pal, J., & Gonawela, A. (2017). Studying political communication on Twitter: the case for small data. *Current Opinion in Behavioral Sciences*, *18*, 97–102.

Pandarachalil, R., Sendhilkumar, S., & Mahalakshmi, G. S. (2015). Twitter sentiment analysis for large-scale data: an unsupervised approach. *Cognitive Computation*, *7*(2), 254–262.

Papacharissi, Z. (2015). *Affective publics: Sentiment, technology, and politics*. Oxford University Press.

Peixoto, T., & Fox, J. (2016). *When does ICT-enabled citizen voice lead to government responsiveness?*

Sapountzi, A., & Psannis, K. E. (2020). Big data preprocessing: an application on online social networks. *Principles of Data Science*, 49–78.

Saul, J. E., Willis, C. D., Bitz, J., & Best, A. (2013). A time-responsive tool for informing policy making: rapid realist review. *Implementation Science*, *8*, 1–15.

Sciences, N. A. of, Medicine, Engineering, D. on, Sciences, P., Science, C., Board, T., Alert, C. on the F. of E., Systems, W., & Directions, R. (2018). *Emergency alert and warning systems: Current knowledge and future research directions*.

Steinkamp, J. M., Pomeranz, T., Adleberg, J., Kahn Jr, C. E., & Cook, T. S. (2020). Evaluation of automated public de-identification tools on a corpus of radiology reports. *Radiology: Artificial Intelligence*, *2*(6), e190137.

Stimson, J. (2018). *Public opinion in America: Moods, cycles, and swings*. Routledge.

Tang, D., Qin, B., & Liu, T. (2015). Deep learning for sentiment analysis: successful approaches and future challenges. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, *5*(6), 292–303.

Torfi, A., Shirvani, R. A., Keneshloo, Y., Tavaf, N., & Fox, E. A. (2020). Natural language processing advancements by deep learning: A survey. *ArXiv Preprint ArXiv:2003.01200*.

Valle-Cruz, D., Criado, J. I., Sandoval-Almazán, R., & Ruvalcaba-Gomez, E. A. (2020). Assessing the public policy-cycle framework in the age of artificial intelligence: From agenda-setting to policy evaluation. *Government Information Quarterly*, *37*(4), 101509.

Yacouby, R., & Axman, D. (2020). Probabilistic extension of precision, recall, and f1 score for more thorough evaluation of classification models. *Proceedings of the First Workshop on Evaluation and Comparison of*

NLP Systems, 79–91.
Yadav, A., & Vishwakarma, D. K. (2020). Sentiment analysis using deep learning architectures: a review. *Artificial Intelligence Review*, *53*(6), 4335–4385.