

Comparative Sentiment Analysis on News Coverage of AI Risks and Regulation using Rule-based and Transformer-based Models

Marwan Noor Fauzy¹, Deni Kurnianto Nugroho², Kardilah Rohmat Hidayat³

Department of Information System, Faculty of Computer Science

Universitas Amikom Yogyakarta

Yogyakarta, Indonesia

marwannoorfauzy@amikom.ac.id¹, deni@amikom.ac.id², kardilah.rh@amikom.ac.id³

Abstract—Rapid development of artificial intelligence technology has raised concerns regarding ethical risks, governance, and the need for adequate regulation. This study aims to analyze the dynamics of public opinion through media coverage of AI risks and regulation. A total of 2,126 news articles were collected from November 1, 2022, to July 28, 2025, using Google News RSS feeds filtered for five major international outlets: Reuters, Bloomberg, The Guardian, CNBC, and The New York Times. The analysis pipeline included article extraction, text cleaning, sentiment classification, and visualization of trends and distributions. Two sentiment analysis approaches were employed: VADER, a widely used rule-based model for news sentiment, and Multilingual BERT (from nlptown), a transformer-based model proven effective in prior studies. Results show that VADER tends to assign neutral labels, while BERT is more sensitive to sentiment nuances. Correlations between the models reveal consistent general trends, but diverge during periods of regulatory activity or ethical controversies. Temporal visualizations indicate spikes in negative sentiment around the enactment of AI regulations. This study concludes that a multi-model approach captures a broader sentiment spectrum. Limitations include restricted media scope, potential data bias, and limited domain-specific sensitivity of the models. Future work should consider expanding media sources, using models trained on AI policy discourse, and integrating entity recognition to identify key actors.

Keywords : artificial intelligence, sentiment analysis, VADER, BERT, news

I. INTRODUCTION

Since the launch of ChatGPT by OpenAI in late 2022, the use of generative Artificial Intelligence (AI) technology has skyrocketed. Media attention on issues such as AI safety, AI risk, AI governance, deepfakes, and regulations such as the AI Act has increased dramatically. This surge has fueled global policymaking and public discourse on responsible AI governance, including initiatives like the UK AI Safety Institute and the EU AI Act [1][2].

The news media has a significant influence on shaping public perception and policy direction. The way news reports on AI risks and regulations influences public opinion, particularly regarding the potential dangers, benefits, and urgency of regulation. To understand this, sentiment analysis of news corpuses is an effective strategy.

Two commonly used sentiment analysis approaches are lexicon-based methods and transformer-based machine learning methods. VADER (Valence Aware Dictionary for sEntiment Reasoning) is an efficient and interpretive rule-based model, particularly suited to short, informal texts [3]. This model has been successfully used in various domains, including predicting election results based on Twitter data, demonstrating that lexicon-based approaches like VADER have relatively high accuracy when applied to specific socio-political contexts [4]. Meanwhile, transformer-based models like BERT offer advantages in understanding complex semantic contexts, often yielding superior performance in formal text classification such as news [5][6].

Saha et al.'s study compared the performance of VADER and BERT in the context of the COVID-19

pandemic on social media platforms and found that BERT consistently delivered higher accuracy in detecting emotional nuances [5]. Furthermore, studies like “Landscape of Generative AI in Global News” by Lu et al. used a RoBERTa-based model for sentiment analysis of a global corpus of news stories about AI, demonstrating that the transformer approach is capable of identifying differences in sentiment between news stories about innovation, security, and regulation [6]. Other studies, such as “Media and Responsible AI Governance”, also emphasize the importance of media as a soft regulatory mechanism in strategic models of interaction between developers, policymakers, and the public [7].

However, the majority of current literature is limited to sentiment analysis in the social media, financial, or political domains. There are still few comparisons between VADER and BERT in the context of news media discussing AI risks and regulation in a structured and longitudinal manner. This research focuses on a comparative analysis of AI news sentiment using keywords such as: AI safety, AI risk, OpenAI, deepfake, chatGPT, AI act, AI governance, for the period from November 1, 2022, to July 28, 2025. This period was chosen because it encompasses the early phase of generative AI's explosive growth and the initial phase of the global regulatory response.

The main contributions of this study are (1) Providing a temporal analysis of AI news sentiment related to risk and regulation throughout 2022–2025, (2) Comparing sentiment classification using two approaches: VADER (rule-based) and BERT (transformer-based), with quantitative evaluations such as accuracy and F1-score, (3) Evaluating the differences in sentiment direction (positive, negative, neutral) of these

two methods and their implications for media framing and public perception, (4) Providing methodological recommendations for researchers and policymakers regarding the most appropriate sentiment analysis model for the regulatory AI domain.

Thus, this study not only provides empirical insights into the framing of AI risk and regulation in media reporting but also underscores the importance of choosing the right sentiment analysis model in studies in this highly dynamic technology domain. This study analyzes a dataset of 2,126 news articles collected from five major international media outlets between 2022 and 2025. By applying both rule-based and transformer-based sentiment analysis models, it aims to uncover patterns of sentiment shifts in relation to key regulatory events and ethical controversies in AI development.

II. RESEARCH METHODS

This study aims to compare rule-based and transformer-based sentiment analysis approaches in understanding public opinion on the risks and regulation of artificial intelligence (AI) through online media coverage. To achieve this goal, a series of processes were carried out, including news data collection, text preprocessing, sentiment analysis, and visualization and evaluation of the results. The complete stages of the method are described as follows:

2.1 Data Collection and Preprocessing

The initial stage of this research involved collecting data from online news articles related to Artificial Intelligence (AI) issues, with a focus on risk and regulation. This process was automated using a Python programming approach based on asynchronous scraping and RSS feed processing. The time period specified in this research was from November 1, 2022, to July 28, 2025, corresponding to the moment when AI issues began to gain widespread global attention. Data was collected from online news articles discussing issues such as AI safety, AI governance, AI regulation, and phenomena like deepfakes. This process was designed to be automated and efficient, encompassing the following steps:

1) Keyword Formulation and Time Range

The data collection was based on a combination of previously formulated relevant keywords, such as "AI safety," "AI governance," "OpenAI," "deepfake," and others. The timeframe was set from November 1, 2022, to July 28, 2025, divided into 25-day intervals to ensure even temporal distribution.

2) Article Search via Google News RSS

Articles were collected by searching Google News RSS feeds based on predetermined keyword combinations. Each query was limited to a maximum of 100 articles to avoid duplication and redundancy. Academic or irrelevant domains such as ieee.org, jstor.org, and sciencedirect.com were excluded from the search results.

3) Link Resolution and Content Extraction

Articles found often lead to redirect links. Therefore, the system automatically resolves redirects to the final URL. Article content is automatically extracted using libraries like `newspaper3k`, which facilitates parsing HTML structures into clear text. For sites requiring dynamic rendering, Selenium with undetected-chromedriver is used to load content as it appears in the browser.

4) Article Validation and Screening

Each extracted article is checked for publication date to ensure it fits within the specified timeframe. Duplicate, empty, or blocked domain articles are ignored and not retained. This validation is crucial to maintaining the quality and cleanliness of the dataset.

5) Logging and Documentation

The entire process is supported by a structured logging system. Information such as the number of successfully retrieved articles, failed processing, and the cause of the failure is recorded in a daily log. This allows the scraping process to automatically resume from its last point if the system is stopped or restarted.

6) Parallelization of Execution for Time Efficiency

To speed up the data collection process, the system runs in parallel using an asynchronous and multi-threaded approach. This allows the extraction of articles from multiple sources and dates to be performed simultaneously without blocking the main execution.

After filtering, deduplication, and validation, the final dataset consisted of 2,126 unique news articles published between November 1, 2022, and July 28, 2025. Each article record includes metadata such as title, publication date, and news source, along with the original and cleaned content. This structured dataset was stored in both JSON and CSV formats, enabling reproducibility and facilitating further sentiment analysis.

2.2 Sentiment Analysis

At this stage, a sentiment analysis process is conducted to classify opinions contained in news content related to the topic of AI and its regulation. The approach used involves a combination of two methods: a lexicon-based model and a transformer-based model. The purpose of using these two approaches is to compare the accuracy, sensitivity, and classification tendencies of each model on the same data, as has been applied in various similar comparative studies [8].

The lexicon-based method was chosen because it is lightweight, transparent, and does not require retraining on new data. One method used is VADER (Valence Aware Dictionary and sEntiment Reasoner), which was specifically developed for analyzing short texts in English and has proven effective in the domains of social media and news [3]. VADER calculates a polarity score from text and categorizes sentiment into three classes: positive, negative, and neutral.

This model is known for its stable accuracy in detecting explicit emotions with minimal computational requirements, making it highly suitable for the baseline analysis in this study.

On the other hand, transformer-based approaches like BERT offer advantages in understanding linguistic context and semantic relationships between sentences. For this purpose, we used a multilingual model from Hugging Face, namely `nlptown/bert-base-multilingual-uncased-sentiment`, which has been trained on multilingual data and is able to provide sentiment ratings on a scale of 1–5. These scores are then mapped into three sentiment categories: positive (score 4–5), neutral (score 3), and negative (score 1–2). This model was chosen because of its ability to capture contextual nuances in long texts such as news articles, as well as its support for multiple languages, which is crucial in the analysis of international news sources [9].

1) Sentiment Model Preparation

VADER is used to obtain a polarity score (compound score) for each article, which is then categorized into positive, negative, or neutral sentiment. Meanwhile, Multilingual BERT is used to predict sentiment based on the context of longer and more complex texts. This model provides a probability distribution for each score, which is then mapped to a sentiment label.

2) Sentiment Analysis Process

Each article was analyzed using both approaches. The VADER model calculated a score based on the intensity of emotionally charged words, while the BERT model generated predictions based on vector representations of entire sentences. Article data was sourced from .json files containing news content, publication dates, and sources. Preprocessing was performed uniformly to remove special characters, excess whitespace, and ensure consistency of input formatting. Both models then analyzed the text and provided sentiment predictions along with confidence scores.

The results from both models were compared to identify similarities, classification differences, and potential biases of each approach. This step is crucial in assessing the reliability of the results, especially for topics with potentially ambiguous or complex framing, such as AI regulatory issues.

3) Saving Results and Output

All analysis results are stored in CSV format, including: file name, publication date, source, sentiment predictions from VADER and BERT, confidence scores, and an initial summary of the article's content. This file then serves as the basis for visual analyses such as time-series plots, weekly comparisons, and mapping sentiment trends against AI policy dynamics.

2.3 Sentiment Visualization

This stage aims to illustrate the temporal dynamics of public sentiment on the topic of AI and its regulation, based on the results of the previously obtained sentiment analysis.

The visualization is performed by combining sentiment data from two models (VADER and BERT) over a specific time period, making it easier to observe patterns and changes in sentiment. The steps taken in this stage include:

1) Data Reading and Aggregation

The sentiment analysis data, stored in CSV format, is read and processed using Python libraries such as pandas for data manipulation. The date columns in the data are then converted to a periodic format (daily, weekly, or monthly) as needed for the analysis. For each time period, the number of articles is aggregated by sentiment category (positive, neutral, negative), separately for each model.

2) Summary Dataset Creation

To facilitate further statistical analysis, a summary dataset containing the total number of articles per sentiment and period was also created. This file was saved in CSV format to document the aggregation results.

3) Data Visualization

Sentiment visualization was performed using a line chart to display sentiment trends over time. This graph was constructed using the matplotlib library and depicts fluctuations in the number of articles per sentiment category, making it easier to observe changes in public perception of AI issues longitudinally. This visualization is differentiated by model (VADER and BERT) to facilitate comparative analysis. Colors are used consistently for each sentiment category: green for positive, blue for neutral, and red for negative.

4) Saving Output

The resulting trend graph is saved as a high-resolution image (.PNG) using functions from the matplotlib.pyplot library. This file is then used in the discussion section to support the interpretation of public opinion dynamics recorded in online media.

2.4 Initial Interpretation and Validation

This stage aims to provide an initial understanding of the previously visualized sentiment analysis results and to conduct initial validation of the accuracy and consistency of the model used. This process serves as a bridge to more in-depth discussions in the next chapter.

1) Trend Pattern Interpretation

Once sentiment visualization data is obtained, the first step is to examine emerging trends over a specific time period. For example, if there is a spike in negative sentiment toward artificial intelligence (AI) issues during a specific period, then an investigation is conducted into the context or accompanying events, such as government policy announcements, publications of controversial research results, or incidents of ethical violations by AI systems.

2) Comparison of Results Between Models

The sentiment analysis results obtained from the two VADER and BERT models were directly compared to assess

their consistency in capturing public opinion dynamics. Although they have different approaches (rule-based vs. transformer-based), similar general patterns can strengthen confidence in the validity of the trends found. Conversely, if significant differences are found, an investigation is conducted to determine possible causes, such as differences in the models' sensitivity to sentence context or the dominance of certain technical terms in articles.

3) Initial Validation

As an initial validation step, a manual inspection of article samples from each sentiment category was conducted. The goal was to ensure that the sentiment labels assigned by the model accurately matched the news content, contextually. This validation is crucial to avoid misleading interpretations, especially given the potential for noise or ambiguity in news texts. The validation results were also used to evaluate the relative performance of the two models in the context of English-language news data discussing AI and its regulation.

2.5 Sentiment Model Evaluation

To ensure the reliability of the sentiment analysis approach used, an evaluation was conducted on the two main methods implemented in this study: VADER (Valence Aware Dictionary and sEntiment Reasoner) and BERT Multilingual. The evaluation focused on the suitability of the classification to the sentence context, sensitivity to nuances of opinion, and consistency of results across time and news sources.

1) Sentiment Output Comparison

The initial evaluation was conducted by comparing the output of the two models across several dimensions: (1) aggregate distribution of sentiment labels (positive, neutral, negative); (2) visualization of sentiment trends in time series for each model; and (3) correlation between the classification results of the two models to measure the level of agreement between the approaches. This analysis aimed to identify whether the two methods produced similar classification trends or whether there were significant differences in sentiment assessments for the same news corpus.

2) Sentiment Evaluation Based on AI Issue Topics

Rather than evaluating individual articles, this approach examines the model's performance in classifying sentiment across key topics emerging in public narratives related to AI risks and regulation. The evaluation focuses on the model's sensitivity in capturing opinions on issues such as: (1) AI regulation and policy (2) Algorithmic bias and fairness (3) Job disruption and automation (4) Privacy and surveillance (5) Existential risks and misuse of AI.

Using automated topic classification methods (e.g., keyword mapping or topic models like BERTopic), the sentiment distribution per topic is analyzed. The results are evaluated to determine whether the model is able to recognize a higher level of negativity in topics such as surveillance or job cuts due to AI, compared to topics like regulation or ethics. This approach provides a macro view of

the model's ability to capture opinions across diverse thematic contexts and avoids selection bias that can arise from using only a select sample of articles.

3) Accuracy and Efficiency Considerations

Model selection also took into account processing efficiency and contextual accuracy. The VADER model was chosen due to its advantages in speed and efficiency for large-scale analysis. However, this model is relatively less sensitive to complex linguistic variations. In contrast, the BERT model offers higher accuracy in understanding semantic context but requires more computational time and resources. Therefore, in the main implementation of trend analysis, the VADER model is used as the primary model, while BERT is used as a cross-validation tool to ensure the accuracy of the results.

4) Implications for the Validity

The difference in classification results between the two models is an important consideration in the validity of this study's findings. The use of the VADER model primarily allows for efficient trend mapping, while cross-validation using BERT ensures the results are not biased by the limitations of a single approach. This combined approach is expected to increase the credibility of the results' interpretation, particularly in the context of analyzing public opinion dynamics regarding AI risks and regulation.

III. RESULT AND ANALYSIS

This section presents the key findings from a sentiment analysis of news coverage on the risks and regulation of artificial intelligence (AI) using two approaches: rule-based (VADER) and transformer-based (multilingual BERT). The analysis was conducted quantitatively and qualitatively to capture the dynamics of public sentiment toward the issue throughout the research period. The analysis steps included basic dataset statistics, aggregate sentiment analysis, sentiment trends over time, comparison of results between models, and qualitative interpretation of key findings.

3.1 Dataset Statistics

The dataset used in this study consists of a collection of 2,126 news articles from various leading international media outlets, such as Reuters, Bloomberg, The Guardian, CNBC, and The New York Times. These articles were collected using a combination of keywords related to artificial intelligence, AI risks, and AI policy and regulation, covering a period from November 2022 to July 2025.




Figure 1. Distribution of the number of news items per month in the dataset

Figure 1 shows the distribution of the number of articles per month. It can be seen that the volume of news coverage varies significantly from month to month. The highest number of articles was recorded in May 2025, which is suspected to correlate with increased public and government attention to the launch of several global policies on AI oversight, including reactions to the development of large-scale models like GPT-5.

Meanwhile, the number of articles tended to be lower in the early months of January and February 2025, likely due to the lack of major events directly related to AI regulatory issues during that period. A gradual increase began to appear in March and April, indicating an escalation in media attention to this topic toward the middle of the year.

Overall, the dataset reflects a growing trend of media interest in the ethics, risks, and regulation of AI, which aligns with the increasing adoption of this technology across various sectors and increasing pressure on governments to develop adaptive regulatory frameworks. This even distribution of time and the relevant topic focus provide a strong foundation for further sentiment analysis and model comparisons in this study.

3.2 Aggregate Sentiment Analysis




Figure 2. Comparison of VADER and BERT sentiment distributions

Figure 2 compares the aggregate sentiment distribution of all articles analyzed using two different approaches: VADER, a rule-based model, and BERT, a transformer-based model. The analysis results show striking differences in patterns between the two models. VADER classified the majority of articles (84.3%) as positive (1,774 articles), with only 307 articles (14.6%) categorized as negative, and 45 articles (2.1%) as neutral. BERT showed a more balanced distribution, with 949 articles (45.9%) positive, 1,128 articles (54.5%) negative, and 49 articles (2.4%) neutral.

This difference reflects the fundamental characteristics of the two methods. VADER, designed for short texts such as social media opinion pieces, tends to favorably evaluate sentences with a formal or neutral tone containing positive words, even in news contexts that raise concerns or risks. This leads to a positive bias in its classification results. On the other hand, BERT, trained on large datasets and complex

language structures, has a deeper understanding of context. This allows BERT to capture nuances of concern, risk, or criticism of AI technology in news articles, resulting in a significantly higher proportion of negative articles.

This fundamental difference is important to consider in further analysis, particularly when selecting a model that better represents public perception or media opinion trends on AI risks and regulation. These results also provide the basis for further exploration into analyzing temporal sentiment trends and correlations with important policy events.

3.3 Sentiment Trend Analysis




Figure 3. Comparison of VADER and BERT sentiment trend analysis

Figure 3 displays the monthly sentiment trends of articles discussing artificial intelligence (AI), analyzed using two approaches: a lexicon-based model (VADER) and a transformer-based model (BERT). The analysis spans the period from November 2022 to May 2025, with a total of 2,126 articles.

1) BERT: Sentiment Variation and Negative Tone Dominance

The top panel of the graph shows the sentiment classification results using the multilingual BERT model. Overall, this model produces a relatively balanced sentiment distribution, but still shows a tendency toward negative sentiment dominance.

Negative sentiment peaked in January 2023 with 55 negative articles, and increased again in December 2023 with 54 negative articles. This spike likely correlated with growing public concern about the risks of AI, such as the discourse on stricter regulations or warnings from the scientific community and technology practitioners. A similar trend reappeared in February 2024 and May 2024, with 51 and 50 negative articles, respectively, indicating that the issue of AI risks remains a key topic. While not dominant, positive sentiment remains consistent. The highest number of positive articles was reached in August 2024 (44 articles), followed by February 2023 (41 articles), and June 2023 (39 articles). This increase could be attributed to media coverage

of AI advancements in healthcare, education, or industrial efficiency.

Articles classified as neutral by BERT were very few in almost all months, with an average of under 5 articles per month. This suggests that news about AI tends to be written expressively in both optimistic and pessimistic tones and is therefore rarely perceived as neutral by the model. Overall, BERT successfully captures the complex dynamics of media discourse, including emotional spikes related to external events. This pattern suggests that public perception of AI tends to be critical and fluctuating, in line with global and local issues.

2) VADER: Positive Overestimation Tendency

The bottom panel of the graph displays the results of the analysis using the VADER model, which is based on lexical rules and is commonly used for short texts like tweets or comments. Unlike BERT, VADER shows a clear dominance of positive sentiment across all time periods.

Each month, positive articles dominated the classification, significantly outnumbering the other categories. The highest peak occurred in April 2023 with 89 positive articles, followed by March 2023 and July 2023 with 83 and 81 positive articles, respectively. The ratio of positive articles to total articles in these months reached over 70%. The number of negative articles in VADER's classification remained low and stable, ranging from 10 to 16 articles per month. Even during periods where BERT recorded a spike in negative sentiment, VADER showed only a slight increase—for example, December 2023 recorded only 14 negative articles, significantly lower than the 54 recorded by BERT.

Neutral articles remained at a very low level, with barely any fluctuation. Most months recorded only 1 to 3 neutral articles. This trend suggests that VADER has a systematic positive bias in the context of long-form news articles. This could be due to the model's limitations in understanding complex sentence context, irony, or ambiguous narrative constructions, which often appear in AI journalism reports.

3) Comparison and Methodological Implications

A comparison of the two models yields several important findings. BERT is more sensitive to fluctuations in public opinion, as reflected in the sharp variations in the number of negative and positive articles across months. For example, the difference between positive and negative articles can be quite small (e.g., August 2024: 44 positive vs. 47 negative), indicating sensitivity to a balanced narrative. VADER tends to overestimate positive sentiment, even in months with substantial criticism of AI. This makes the model less reliable for long-form, opinion-heavy media such as news or investigative reports. Both models agree that neutral articles are very rare, confirming that AI coverage tends to be written in a polarized manner (positive or negative), possibly due to the controversial and broad-reaching nature of AI issues.

Therefore, the choice of sentiment analysis model should consider the type of text and the complexity of the context.

For policy or social influence analysis, BERT provides more representative results, while VADER is better suited as an initial indicator for short texts or quick exploration.

3.4 Comparison of Results Between Models

To evaluate the consistency of the results between the rule-based and transformer-based approaches, a comparative analysis was conducted on the sentiment scores generated by the VADER and BERT models on the same set of news articles. Figure 4 shows a scatterplot visualization of the sentiment scores from the two models.




Figure 4. Sentiment Score Correlation between VADER and BERT

1) Correlation Analysis

A scatter plot visualization is used to illustrate the relationship between the sentiment scores from the two models. In Figure X, each dot represents a single article, with the sentiment score from VADER displayed on the x-axis and the score from BERT on the y-axis. The calculation results show that the Pearson correlation coefficient between the sentiment scores from the two models is $r = -0.21$, indicating a weak negative correlation. This means that, in general, the two models tend to produce different sentiment interpretations.

2) Observed Differences

This difference is consistent with findings in previous literature. VADER is a lexicon-based model designed for short, informal texts such as social media [3]. While accurate for direct expressions, VADER has limitations in understanding implicit sentiment and complex sentence context, which often appear in news articles.

In contrast, BERT is a transformer model trained on full sentence context [10]. In several studies, such as [11], [12], BERT has demonstrated superior performance in formal text domains due to its ability to capture non-explicit semantic relations. This causes BERT to often detect polarity (positive/negative) in sentences that VADER would consider neutral.

For example, in a news story addressing the risks of AI regulation to innovation, BERT might classify it as negative towards the tech sector, while VADER would remain neutral due to the lack of explicit negative words.

3) Implications for Analysis

This low correlation confirms that model selection significantly impacts sentiment analysis results, especially in the context of news articles with their rich context and complexity. VADER offers fast and transparent interpretation, while BERT excels at understanding implicit meaning. Therefore, combining both approaches for example, by weighting their results or using them for cross-validation can enrich interpretation and improve the reliability of the analysis.

3.6 Qualitative Analysis

A qualitative analysis of sentiment results reveals several important patterns in media coverage of the risks and regulation of artificial intelligence (AI). Combining quantitative approaches from the VADER and BERT models with observations of news content, we uncover dynamics in public opinion and media attention that reflect growing concerns about AI governance.

1) Growing Concerns Over AI Governance

From January to April 2025, there was a gradual increase in the number of articles containing negative sentiment, particularly related to governance issues, algorithmic bias, and cross-border regulation. Content analysis showed that public concern tended to increase along with discussions surrounding AI's disruption to democracy, privacy, and algorithmic transparency. Negative sentiment peaked in the first week of May 2025, coinciding with the publication of the European Union's draft AI law, which sparked controversy in various global media outlets.

These findings indicate that large-scale policy events have a significant impact on the tone of news coverage, particularly when accompanied by narratives of threat or uncertainty about the future of technology [13].

2) The Spike in Negative Sentiment and the Emergence of the Term Risk

Keywords such as surveillance, deepfake, and autonomous weapons frequently appear in articles with high negative sentiment according to the BERT model. VADER tends to assign neutral labels to articles that use technical terms without explicit emotion, but BERT captures the nuances of implicit concern, reinforcing the finding that context-based approaches are more sensitive to risk framing in technology discourse [12].

3) Dominant Actors and Institutions

Certain institutions and actors appear repeatedly in the corpus of articles. OpenAI, Google DeepMind, and MetaAI are frequently mentioned in the context of innovation and the development of advanced AI models. Sentiment toward them

tends to be polarized, with praise for technical advances but also criticism regarding AI ethics. The European Union and the White House Office of Science and Technology Policy (OSTP) emerge as key regulatory figures, often associated with AI regulatory frameworks and responses to the AI Act. Individuals such as Elon Musk, Sam Altman, and Geoffrey Hinton are frequently mentioned, particularly in discussions about the social impacts of AI and calls for moratoriums on certain technologies.

This finding is consistent with previous studies showing that media framing is highly dependent on an actor's position within the AI ecosystem, whether as an innovator, critic, or regulator [16].

4) Qualitative Conclusion

Overall, media coverage reveals a complex emotional dynamic surrounding AI developments. Negative sentiment stems not always from technological failures, but rather from uncertainty about governance and long-term impacts. A combined approach of sentiment analysis and observations of key actors/institutions provides a richer picture of how AI risk issues are perceived and communicated to the public.

VI. CONCLUSION

This study examines the dynamics of public opinion in media coverage of the risks and regulation of artificial intelligence (AI) through a comparative sentiment analysis approach. The two main models used, VADER (rule-based) and BERT Multilingual (transformer-based), exhibit different characteristics and sensitivity in interpreting news text sentiment.

In general, the BERT model demonstrates greater sensitivity to sentence context and is able to capture implicit nuances, particularly in speculative issues or those containing non-explicit concerns. In contrast, VADER is more conservative, often labeling text with a formal or technical tone neutral. The correlation between the models is moderate to strong, but some significant differences emerge during periods of significant events, such as AI policy announcements or controversial technology launches.

The qualitative analysis reveals that spikes in negative sentiment often correlate with discourse about mass surveillance, algorithmic bias, and the potential misuse of AI technology. Institutions such as the European Union and OpenAI, as well as figures like Elon Musk and Sam Altman, play a significant role in shaping the public narrative around AI governance.

The primary contribution of this research is to provide a more holistic understanding of public sentiment by combining two complementary modeling approaches and highlighting the socio-political context underlying opinion shifts. These findings are expected to support the development of AI policies that are more responsive to public concerns and ethically inclusive. Future extensions of this

study may include multi-modal analysis (e.g., images or videos), the adoption of more advanced LLMs, and the integration of social media data to capture broader and real-time public perceptions.

This research underscores the value of comparative sentiment analysis in informing more nuanced, inclusive, and context-aware approaches to AI governance, offering a foundation for future interdisciplinary exploration.

REFERENCES

- [1] A. Reuel and T. A. Undheim, *Generative AI Needs Adaptive Governance*, arXiv preprint, 2024.
- [2] A. Batool, D. Zowghi, and M. Bano, *Responsible AI Governance: A Systematic Literature Review*, arXiv preprint, 2023.
- [3] C. J. Hutto and E. Gilbert, “VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text,” *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 8, no. 1, pp. 216–225, 2014.
- [4] D. K. Nugroho, “Sentiment Analysis for Predicting the 2020 US Presidential Election Using VADER,” *Proceedings of the 11th International Conference on Cloud Computing, Data Science & Engineering (Confluence)*, Noida, India, pp. 136–141, 2021.
- [5] P. Saha, S. Mahato, and A. Das, “VADER vs. BERT: A Comparative Performance Analysis for Sentiment on Coronavirus Outbreak,” in *Recent Advances in Computational Intelligence*, Singapore: Springer, pp. 407–418, 2023.
- [6] X. Lu, L. Li, Y. Xu, B. Z. Zhang, and L. Hemphill, *Landscape of Generative AI in Global News: Topics, Sentiments, and Spatiotemporal Analysis*, arXiv preprint, 2024.
- [7] N. Balabanova, A. Bashir, et al., *Media and Responsible AI Governance: A Game-Theoretic and LLM Analysis*, arXiv preprint, 2025.
- [8] F. N. Ribeiro, M. Araújo, P. Gonçalves, M. A. Gonçalves, and F. Benevenuto, “SentiBench: A Benchmark Comparison of State-of-the-Practice Sentiment Analysis Methods,” *EPJ Data Science*, vol. 5, no. 1, p. 23, 2016.
- [9] L. Zhang, S. Wang, and B. Liu, “Deep Learning for Sentiment Analysis: A Survey,” *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 8, no. 4, e1253, 2018.
- [10] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” *Proceedings of NAACL-HLT*, 2019.
- [11] W. Yin, K. Kann, M. Yu, and H. Schütze, “Comparative Analysis of BERT-based Models for Sentiment Classification,” *Proceedings of ACL*, 2021.
- [12] T. Hoang, T. H. Nguyen, and M. T. Nguyen, “Benchmarking Sentiment Analysis Models in Financial and Political News,” *Journal of Computational Social Science*, vol. 5, pp. 1–20, 2022.
- [13] A. Jobin, M. Ienca, and E. Vayena, “The Global Landscape of AI Ethics Guidelines,” *Nature Machine Intelligence*, vol. 1, no. 9, pp. 389–399, 2019.
- [14] B. Binns, “Fairness in Machine Learning: Lessons from Political Philosophy,” *Proceedings of the 2020 ACM Conference on Fairness, Accountability, and Transparency*, pp. 149–159, 2020.