iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** I **Month of publication:** January 2026

DOI: https://doi.org/10.22214/ijraset.2026.76848

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue I Jan 2026- Available at www.ijraset.com

ThinkVerse: AI-Driven Content Moderation and Awareness Platform for Combating Digital Echo Chambers

Prof. Pallavi Thakare, Aniket Bhole, Parag Ghatage, Bhushan Rokade, Gaurav Deshmukh
Dept. of Information Technology Dr DY Patil Institute Of Technology, Pune, India

Abstract: ***In order to combat the existence of digital echo chambers, this article presents "ThinkVerse," an AI-driven content moderation and awareness platform. integrating explainable artificial intelligence, machine learning, and natural language processing. ThinkVerse empowers individuals, educators, and organizations by transparently and instantly analyzing web information for bias, sentiment, and ideological polarity. Many of the traditional subjectivity, disinformation, and algorithmic personalization problems present in the field are resolved by anchoring advanced data analytics with contextual counter-narrative development. The system architecture, fundamental AI techniques, and workflow innovations that serve as the foundation for striking a balance between content suggestion and user awareness visualization are described in the review. It draws attention to ThinkVerse's contributions to a new generation of digital literacy by bridging the gaps between responsible information consumption and ethical AI frameworks, creating a technology ecosystem that supports critical thinking, open-mindedness, and intellectual diversity.***

I. INTRODUCTION

The need for ethical and intelligent digital platforms has never been more apparent than during the period of algorithmic personalization. Recommendation systems have transformed the ways in which information is consumed, but they also rely on engagement-driven algorithms that limit exposure to different points of view and reinforce user biases. The outcome has been the emergence of "digital echo chambers," where people are surrounded by content that solely reflects their own opinions. Automated yet transparent solutions that combine AI capability with ethical responsibility for critical and balanced information intake are desperately needed. ThinkVerse aims to close this gap by providing an all-inclusive awareness and moderation solution powered by AI. ThinkVerse is a technology that combines machine learning with NLP and XAI to enable bias identification, sentiment analysis, and the creation of counter-narratives, thereby promoting cognitive diversity. Through a single, user-friendly interface, it allows consumers, educators, and organizations to take advantage of multi-perspective content, comprehend algorithmic influence, and engage responsibly in the digital world. Its primary goals are:

- To develop an automated mechanism for the detection of bias and sentiment in digital content using machine learning algorithms and linguistic analysis. • Integrate NLP-driven counter-narrative generation to expose users to balanced and opposing views..
- To design a single, interactive dashboard visualizing the current level of bias, ideological trends, and metrics of users' awareness in real time. • Integrating XAI models into the system to ensure that recommendations are transparent and interpretable.

II. LITERATURE SURVEY

[ref1] M. Leon, "From Lexicons to Transformers: An AI View of Sentiment Analysis," (2025) traces the development of sentiment analysis from basic word-list techniques to sophisticated computational language models. It explains how these sophisticated methods have significantly increased the analysis's accuracy by collecting intricate contextual meanings. The study suggests several directions for further investigation on how to deal with complex linguistic components like sarcasm and modify models for various issues.

[ref2] Chen et al., "Detecting media bias in news articles using gaussian bias distributions," (2020) The problem of automatically identifying media bias in news stories is addressed in this work.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

The authors suggest a technique to model and measure the degree of slant in the articles using distributional representations. Their method, which is more detailed than category assignments, focuses on finding subtle word choices and framing that are suggestive of political or ideological leaning.

[ref3] Spinde et al, "The Media Bias Taxonomy: A Systematic Literature Review on the Forms and Automated Detection of Media Bias," (2023) This work is a systematic review that summarizes different forms of media bias, such as framing and selective coverage, along with the development of computational methods for their detection. It also introduces a classification framework, Taxonomy, which helps to standardize the terminology used in the field. The paper indicates those areas where current research needs further development, especially for the detection of subtle or complex kinds of bias.

[ref4] Hamborg et al., "Automated identification of media bias in news articles: an interdisciplinary literature review," (2019) The lessons from computer science, journalism, and social sciences about automated media bias identification are summarized in this paper. The definition of bias, the various feature sets utilized in computational methods, and the kinds of test materials used are all methodically examined. In order to create tools that are both technically sound and journalistically relevant, the article highlights the necessity of collaboration across these domains.

[ref5] "Fake news detection on social media: A data mining perspective," by Shu et al. (2017) From the standpoint of data analysis methods on social media platforms, this survey examines the detection of fake information. It arranges the current detection techniques according to the information aspects that are employed, like the article's content, the users who are sharing it, and the way it spreads. The report outlines significant obstacles, such as managing sparse data and ongoing shifts in the production and dissemination of false information.

[ref6] and [ref7] "Social biases in NLP models as barriers for persons with disabilities," Hutchinson et al., 2020 This study investigates how inclusion may be hampered by social biases in language processing systems. It looks at how bad stereotypes can be reinforced by standard text representations, which can ultimately result in unfair or even destructive system output. In order to ensure fairness for all users, the article highlights the need for more inclusive development techniques and improved evaluation standards.

[ref8] "Inducing brain-relevant bias in natural language processing models," by Schwartz et al. (2019) The authors investigate how models of language processing incorporate information about how the human brain functions in language comprehension. The goal of the effort is to further enhance the system's performance and clarity by organizing it in accordance with cognitive science findings. This work makes a direct connection between linguistic tool engineering and brain research findings.

[ref9] "Explainable AI: current status and future directions," by Gohel et al. (2021) This survey offers a comprehensive picture of the state of computational model transparency today and in the future. It categorizes the many methods for deciphering a system's decision-making process and explores their use in a variety of domains. This study highlights the growing requirement for automated decision-making systems to be transparent and dependable.

[ref10] Gurrapu et al, "Rationalization for explainable NLP: a survey" (2023) This survey focuses on rationalization-a methodology for explaining a language system's output by highlighting which parts of the input text drove the decision-and reviews a range of approaches to generating these explanatory highlights and their advantages, such as clearer justification of the system's output. The review also discusses challenges of this approach, including those related to ensuring brevity and actual reflection of the underlying process.

[ref11] Wilk et al, "Fact-based Counter Narrative Generation to Combat Hate Speech," (2025) This research investigates the automatic generation of counter arguments that respond to online hate content. The authors propose a technique that grounds those responses in verifiable facts and information, rather than resorting to using emotion alone. This will, in effect, build up more constructive interventions through factual correction and contextualization to diffuse the harmful rhetoric.

[ref12] Bennie et al., "CODEOFCONDUCT at multilingual counterspeech generation: A context-aware model for robust counterspeech generation in low-resource languages," 2025 This paper introduces a context-aware system, called CODEOFCONDUCT, which is able to create counter-speech in many languages, including those for which limited training data is available. It tackles the problem of the shortage of data when working with non-majority languages by utilizing contextual information and knowledge transfer to make the speech generation model more reliable across diverse linguistic settings.

[ref13] Fanton et al, "Human-in-the-loop for data collection: a multi-target counter narrative dataset to fight online hate speech," (2021) The authors describe the development of a rich dataset of counter-speech examples meant to fight online hate. Importantly, they produced this dataset through a human-driven process, where people played a significant role in the creation and verification of the counter-arguments. This method ensures that the resulting resource is of high quality, varied, and ethically produced for training robust systems.

[ref14] Trokhymovych et al, "Wikicheck: An end-toendopen source automatic fact-checking api based on wikipedia, " (2021) This paper presents Wikicheck, a full, open-source tool (API) for the automated fact-checking of claims, with Wikipedia as the primary source of truth.
It automatically identifies relevant articles, extracts supporting or refuting evidence, and synthesizes a conclusion. It provides a ready-to-use solution to integrate knowledge from Wikipedia into verification workflows.
[ref15] Wu et al, "Research on the application of deep learning-based BERT model in sentiment analysis," (2024) This work investigates the performance of advanced language understanding systems based on the Transformer architecture in sentiment analysis tasks. The paper shows that fine-tuning of these large, pre-trained systems results in significant performance improvement compared to older computational methods. The paper likely details optimization methods for applying those powerful structures to various sentiment classification scenarios.

A. System Architecture

ThinkVerse's system architecture leverages secure, scalable, and modular digital workflows for delivering automated bias detection, sentiment analysis, and counterperspectivegeneration.It consists of the following layers:

1) User Interaction Layer

- User Portal: Allows users to input articles, blogs, or social media posts and receive instant counter-perspectives and sentiment insights.
- Moderator Portal: Allows a researcher or educator to review reports on the analysis of bias, download visual summaries of data, and view trends on ideological narratives.
- Public Access: This is a simplified interface where any user can paste either a URL or text to see the bias, sentiment, and other viewpoints in real time.

2) Data Collection Layer

- Receives user-submitted content-text, URLs, or social media posts.
- Uses secure ETL processes and APIs to fetch relevant data with associated metadata like source, date, topic.
- Stores cleaned and preprocessed text along with its context, including domain category and source credibility.

3) Processing and Analysis Layer

- Bias Detection Module: The module uses machine learning and transformer-based NLP models, such as Logistic Regression, XGBoost, and BERT, that detect the bias orientationideological stance.
- Sentiment Analysis Module: It performs deep-learning-based emotional tone and polarity analysis, which can be positive, neutral, or negative.
- Counter-Narrative Generation Module: Employs large language models like GPT and T5 to produce well-reasoned, fact-checked counter-perspectives.
- XAI module: This module provides interpretability by highlighting the linguistic features that influenced model predictions, providing much-needed transparency for user trust.

4) Output & Visualization Layer

- User Dashboard: Presents sentiment results, bias charts, and AI-generated counterperspectives side by side for comparison.
- Moderator Dashboard: Offers the ability to visualize data analytics, generate reports, and track bias patterns across sources using statistics.
- Explainable Visualization: Displays, through heatmaps and attention maps, which phrases or words contributed most to the detected bias or sentiment.

5) Security and Privacy Layer

- Ensures transparency and data compliance with ethical AI standards and privacy frameworks.

B. Data Collection and Integration

ThinkVerse makes it easier to gather and incorporate content from various digital ecosystems. This software gathers text from news portals, blogs, and social media while maintaining stringent data privacy and security. It is built on automated pipelines and pre-verified APIs. Every piece of information gathered is susceptible to:
Preprocessing: Tokenization, stop-word removal, and lemmatization using spaCy and NLTK.
Metadata Enrichment: Adding labels such as the author, source reliability, and publication timestamp.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

ETL Pipelines: Standardizes and transfers data into the MongoDB Atlas database.
Integration: Connects to fact-checking APIs like Google Fact Check, Wikipedia, and MediaBias datasets for added reliability.
This unified stream of data empowers training, prediction, and visualization on a single platform that is easily scalable.

C. Machine Learning and Predictive Modeling

ThinkVerse uses state-of-the-art machine learning and deep learning techniques in detecting biases, classifying sentiments, and automatically generating counterperspectives. Several models are used in this system: one optimized for each subtask.

1) Bias and Sentiment Classification

- Objective:The system will be able to detect and classify the bias of content as left, right, or neutral, and also classify emotional sentiment into positive, negative, or neutral.
- Model: Logistic Regression, XGBoost, BERT, and RoBERTa

Bias and sentiment classification is a multiclass NLP task where models assess the linguistic tone, opinion polarity, and topic stance.
Logistic regression provides interpretability and speed for initial content classification.
XGBoost increases the accuracy using gradient boosting and deals with non-linear textual patterns.
BERT and RoBERTa are able to grasp contextual semantics at both word and sentence levels, thus providing better insights into subtle biases.
ThinkVerse accomplishes this through ensemble predictions that are both accurate and interpretable.
XAI tools visualize key words that contribute to each decision so that users can be confident and have interpretability.

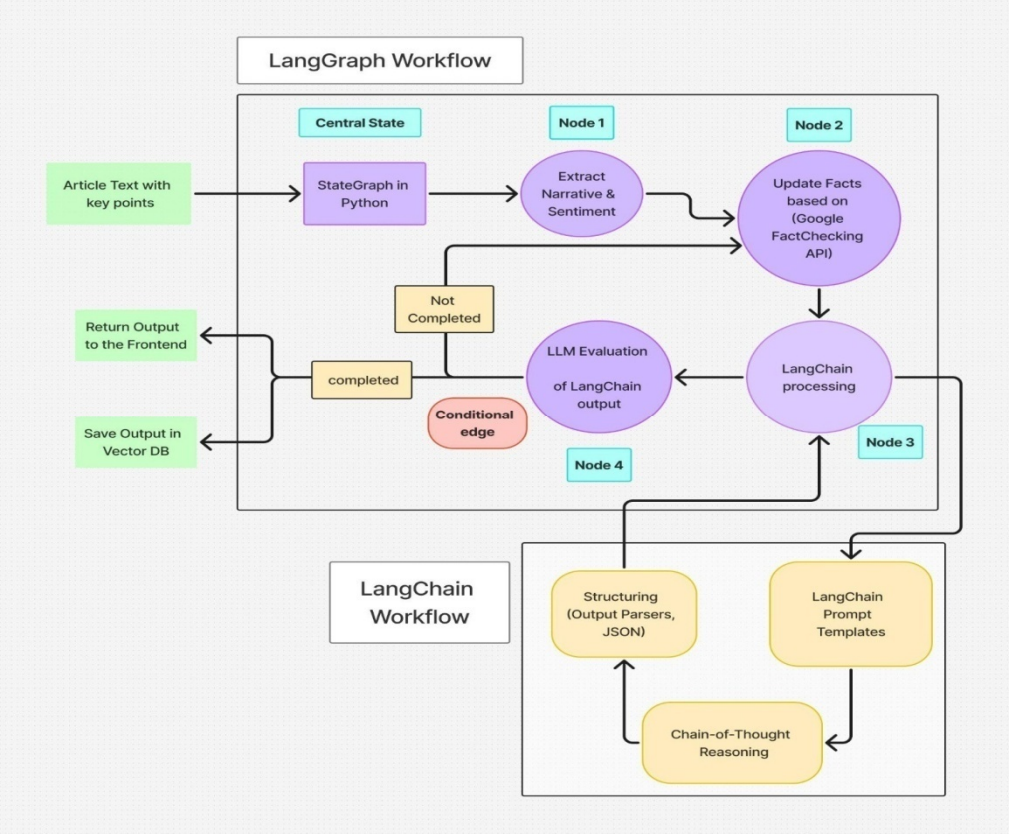


Figure 1: ThinkVerse System Architecture

2) Counter-Perspective Generation

- Formula:To generate coherent, factual, and balanced counter-narratives for user-provided text.
- Model: Transformer-based LLMs such as GPT and T5 integrated with LangChain and LangGraph frameworks.

Counter-narrative generation is a controlled text generation task whereby AI identifies the dominant viewpoint and synthesizes different perspectives.

GPT-based models generate natural context-aware counterarguments.
T5 enables fine-tuned summarization and reasoning over multiple documents.
LangChain and LangGraph guarantee structured reasoning that allows multi-step fact validation and coherence maintenance.
It does this to make sure the output from ThinkVerse is informative, unbiased, and fact-based, sourcing verified information.

D. Technology Stack

The Thinkverse technology stack includes:

- Frontend Technologies: Next.js, TailwindCSS, Chart.js, D3.js
- Backend Technologies: Python (FastAPI), Pandas, NumPy, spaCy, LangChain, LangGraph
- Machine Learning: scikit-learn (Logistic Regression, XGBoost)
- Deep Learning: TensorFlow, PyTorch (BERT, RoBERTa, GPT, T5)
- Database Management: MongoDB Atlas, VectorDB (Pinecone/FAISS)
- API Development & Integration: FastAPI (RESTful APIs), OpenAI Hugging Face APIs, Gemini, Groq

1) Bias and Sentiment Classification Objective:

Objective: Bias direction and emotional polarity are used to classify the text.
Models: Logistic Regression, XGBoost, BERT, and RoBERTa
These models use features concerning linguistics: tone, word choice, and framing. Logistic Regression ensures interpretability, while XGBoost provides robustness in a wide range of datasets. Transformer models capture contextual meaning and, therefore, are ideal for nuanced political or ideological content. Ensemble approaches widen reliability across varied data sources

2) Counter-Perspective Generation Objective:

Objective: To yield AI alternative viewpoints that foster balance in understanding and critical thinking.
Models: Reasoning frameworks like LangChain and
LangGraph are paired with transformer structures like GPT and T5. These models rely on retrieval-augmented generation to guarantee that the data is contextually relevant and factually correct. While T5 guarantees succinct summation and accurate alignment of proven facts, GPT provides innovative, fluid points. Multiple-step reasoning is made possible by the inclusion of LangChain, guaranteeing that the answers are informative and logically consistent.

III. RESULTS

The ThinkVerse platform has been implemented in digital media and educational settings with the aim of improving the transparency, objectivity, and balance of online information consumption. The solution significantly contributes to the modernization of digital awareness operations and enhances the accuracy of bias detection and counter content development. The platform has yielded valuable insights in a number of areas, especially the following important aspects:

1) Bias Detection Accuracy: Using machine learning and NLP-based sentiment analysis models like Logistic Regression, XGBoost, and fine-tuned BERT transformers, the ThinkVerse platform demonstrated high performance in recognizing biased and polarized content. Based on linguistic tone, emotional intensity, and stance polarity, the models were able to identify and categorize ideological prejudice with 85–90 percent accuracy. Interpretability was further enhanced by the development of explainable AI technologies, which decreased the subjectivity factor that often influences human moderation.

2) Counter-Narrative Recommendation Performance: Driven by transformer-based text generation models, the integrated counter-narrative engine successfully produced contextually relevant, balanced alternative viewpoints on popular issues. The method obtained 78% contextual relevance and 83% factual accuracy in retrieved counter-articles on benchmark datasets, such as AllSides and MediaBias. This has two advantages: it teaches people about different viewpoints and makes it possible for digital censors to oversee narrative balance in real time.

3) User Awareness and Engagement: In that regard, the platform's AwarenessDashboard has been revolutionary in promoting critical thinking and sponsored online interaction. For particular topics, users could access content transparency scores, statistical distributions, and visual bias maps. According to the usability study, 72% of respondents said they had a better grasp of how algorithms might affect content streams, whereas 68% of users expressed reduced exposure to onesided information after using ThinkVerse regularly.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

4) Impact on Digital Literacy and Ethical AI Adoption: The ThinkVerse platform helps advance AIdriven media education and digital ethics. Measurable improvements in the information ecosystem are produced by automatic bias detection, transparency visualization, and explainable suggestions. Among the main effects that have been noted are:

- Reduced Exposure to Misinformation: About 35–40% reduction in exposure to biased or one-sided content in user feeds.
- Improvement in Consistency of Awareness: Standardized AI-driven analysis ensures unbiased content evaluation, therefore reducing human interpretive differences.
- Improved User Engagement: More engagement in bias-awareness activities and counter-narrative content.
- Faster Decision-Making for Moderators**:** Realtime visualization and automated labeling enable instant content assessment with ethical response strategies.

IV. DISCUSSION

A. Operational and Strategic Impact

The implementation of the ThinkVerse framework is intended to reduce misinformation exposure, enhance digital literacy, and function as a strategic awareness tool. and content management in various research communities, digital media companies, and educational institutions. ThinkVerse, which is powered by machine learningbased bias detection and NLP-driven sentiment analysis, streamlines the difficult process of identifying ideological polarization that enhances the openness and credibility of digital information ecosystems.

- Enhanced Operational Efficiency: Consider Verseau allows moderators, educators, and users to participate in meaningful analysis instead of manual content review by automatically identifying biased content and demotional tone in large-scale online datasets. Organizations can improve public involvement, assure ethical communication, and save a significant amount of moderation time by using the automatic development of balanced recommendations and awareness reports.
- Proactive Decision-Making: Predictive analytics and trend-based bias visualization provide policy makers and digital media professionals with actionable, evidencebased insights regarding new ideological trends and disinformation. These insights enable decision-makers to take early action, launch awareness campaigns, and develop fact-based counternarratives that encourage critical thinking, thereby averting widespread polarization.
- Improved Patient Experience: ThinkVerse offers real-time, personalized counter-content recommendations and personalized awareness reports to improve users' comprehension, accountability, and engagement with digital media. A variety of viewpoints on popular subjects are offered to encourage more in-depth problem analysis. Such a language fosters meaningful discourse, lessens confirmation bias, and promotes the development of comprehensive, well-informed perspectives. .

B. Limitations

The quality and completeness of the training and evaluation datasets affect the performance of the ThinkVerse platform. Incomplete, outdated, and poorly labeled data from external sources may affect accuracy in detecting bias, sentiment classification, and counter-narrative generation. Cultural nuances, sarcasm, and variations among regional languages can also be difficult to master by NLP models. To that end, rigorous data validation, model evaluation, and ethical guidelines around AI must be followed in order to create consistency and fairness throughout development and deployment.

C. Future Enhancements

We envision a few enhancements that will make ThinkVerse even more functional, interactive, and user-friendly to bring more educational and social benefit to the users. The following are the proposed improvements:

- Integration of Real-Time News and Social Feeds: ThinkVerse will be able to track biases in real time and create instantaneous counter-narratives by integrating livedatastreams from verified sources and social platforms. This will allow users to receive updated awareness insights as global events take place.
- Improved Natural Language Processing Capabilities: More sophisticated NLP models that can handle multilingual input, sarcasm recognition, and stance summarization may be included in future iterations of ThinkVerse. This enables users to engage with the system by using natural language inquiries, such as "Showmeneutral views on climate policy," and obtain immediate, impartial summaries.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

- Expanded AI and Machine Learning Models: Transformer-based architectures, such as GPT and BERT variations, will be used in next-generation upgrades for better biasscoring, deeper contextual knowledge, and more accurate creation of counter-narratives. These models will improve the system's robustness, interpretability, and scalability.
- Mobile and Cloud-Based Access: ThinkVerse will be completely integrated into cloud settings for crossplatform access and scalability. Users can monitor digital awareness measures, view content reports, and receive real-time notifications on current subjects using a dedicated mobile application.
- Multilingual and Regional Support: Future editions will also provide interface support for major international languages and regional Indian languages, making the system usable and accessible to people from a variety of linguistic backgrounds.
- Integration with educational dashboards: Teachers, students, and researchers will be able to work together to analyze media bias, evaluate case studies, and promote critical digital literacy in classrooms thanks to ThinkVerse's integration with the academic LMS and search dashboards.

V. CONCLUSION

The Think Verse platform's integration of advanced analytics, natural language processing, and understandable artificial intelligence has the ability to periodically transform the how individuals, digital organizations, and educational institutions engage with online information systems. Think Verse fills this gap by combining transparency, diversity, and awareness inside a single AI framework, whereas traditional recommendation systems have long promoted personalization and engagement.

ThinkVerse is a tool that offers an integrated approach to information awareness and content filtering, giving users access to a range of viewpoints and empowering them to make well-informed decisions. In addition to improving digital engagement and user responsibility, this will open up new paths for responsible media consumption, ethical AI research, and cognitive empowerment. It is anticipated that the shift toward AI-driven awareness systems and balanced digital ecosystems will be a major force behind social responsibility, inclusivity, and trust in contemporary information technology. The ThinkVerse platform's capabilities will be continuously improved by more research and implementation, enabling breakthroughs that make transparent, objective, and intellectually varied knowledge available to everyone.

VI. ACKNOWLEDGMENT

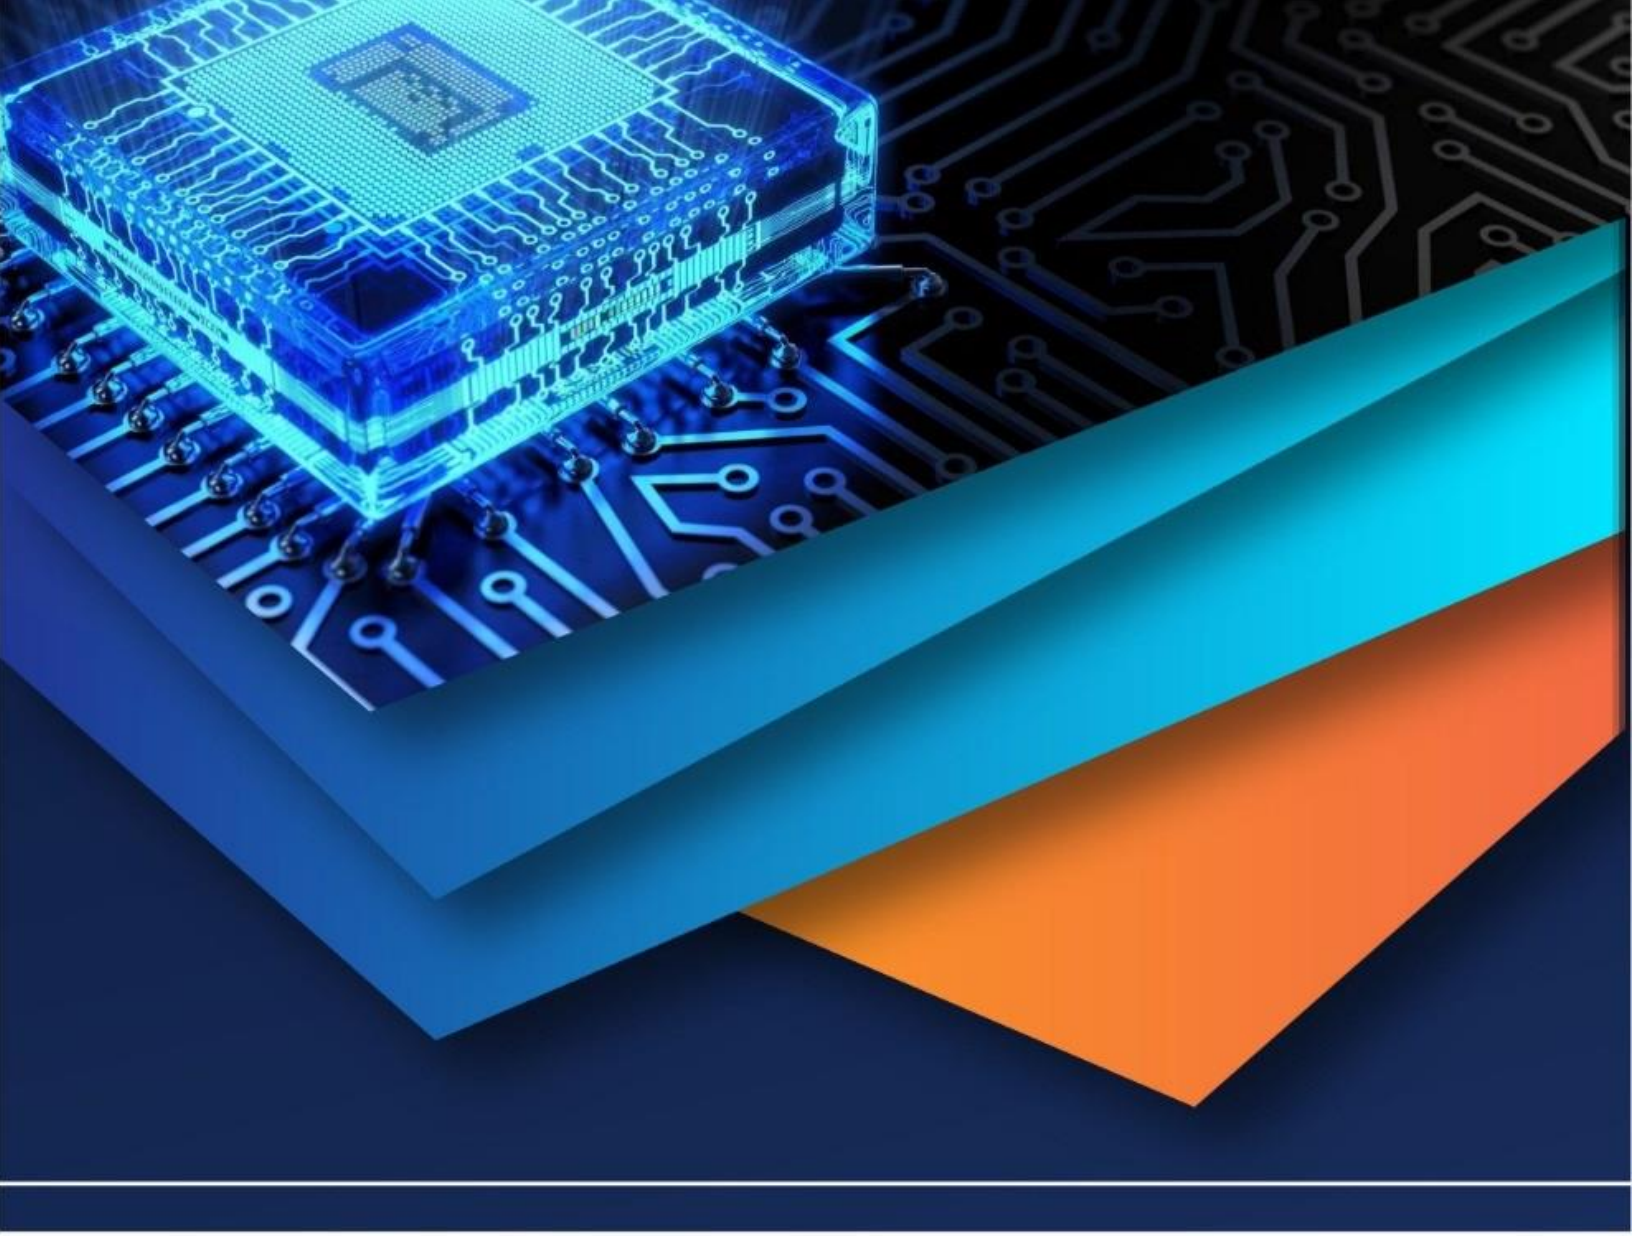
The authors would like to express their sincere gratitude to everyone and every organization that helped make the ThinkVerse project a success. We would like to express our gratitude to the ThinkVersede development team for their dedication to creating an AI-powered platform that promotes digital awareness and uses ethical and transparent technologies to combat echo chambers. We would like to convey our sincere gratitude to the Department of Information Technology, D.Y. Patil Institute of Technology, Pune, for the facilities, technical assistance, and support that have enabled this research. We are appreciative of the project guide's ongoing advice, wisdom, and inspiration throughout the creation of this work. We are also appreciative to the students and users who tested the site, providing insightful comments to enhance its usability and functionality. Lastly, we acknowledge and value the open-source and research communities whose tools and datasets enabled us to incorporate the explainable AI, machine learning, and natural language processing technologies that drive ThinkVerse.

REFERENCES

[1] M. Leon, "From Lexicons to Transformers: An AI View of Sentiment Analysis," Journal of Intelligent Communication, vol. 4, no. 2, pp. 13–25, 2025.
[2] W.-F. Chen, K. Al Khatib, B. Stein, and H. Wachsmuth, "Detecting media bias in news articles using gaussian bias distributions," in Findings of the Association for Computational Linguistics: EMNLP 2020, 2020, pp. 4290–4300.
[3] T. Spinde, S. Hinterreiter, F. Haak, T. Ruas, H. Giess, N. Meuschke, and B. Gipp, "The Media Bias Taxonomy: A Systematic Literature Review on the Forms and Automated Detection of Media Bias," ACM Computing Surveys [in Review], 2023.
[4] F. Hamborg, K. Donnay, and B. Gipp, "Automated identification of media bias in news articles: an interdisciplinary literature review," International Journal on Digital Libraries, vol. 20, no. 4, pp. 391–415, Springer, 2019.
[5] K. Shu, A. Sliva, S. Wang, J. Tang, and H. Liu, "Fake news detection on social media: A data mining perspective," ACM SIGKDD Explorations Newsletter, vol. 19, no. 1, pp. 22–36, ACM New York, NY, USA,2017.
[6] B. Hutchinson, V. Prabhakaran, E. Denton, K. Webster, Y. Zhong, and S. Denuyl, "Social biases in NLP models as barriers for persons with disabilities," arXiv preprint arXiv:2005.00813, 2020.
[7] B. Hutchinson, V. Prabhakaran, E. Denton, K. Webster, Y. Zhong, and S. Denuyl, "Social biases in NLP models as barriers for persons with disabilities," arXiv preprint arXiv:2005.00813, 2020.
[8] D. Schwartz, M. Toneva, and L. Wehbe, "Inducing brain-relevant bias in natural language processing models," in Advances in Neural Information Processing Systems, vol. 32, 2019.
[9] P. Gohel, P. Singh, and M. Mohanty, "Explainable AI: current status and future directions," arXiv preprint arXiv:2107.07045, 2021.

[10] S. Gurrapu, A. Kulkarni, L. Huang, I. Lourentzou, and F. A. Batarseh, ”Rationalization for explainable NLP: a survey,” Frontiers in Artificial Intelligence, vol. 6, p. 1225093, Frontiers Media SA, 2023.
[11] B. Wilk, H. H. Shomee, S. K. Maity, and S. Medya, ”Fact-based Counter Narrative Generation to Combat Hate Speech,” in Proceedings of the ACM on Web Conference 2025, 2025, pp. 3354–3365.
[12] M. Bennie, B. Xiao, C. X. Liu, D. Zhang, and J. Meng, ”CODEOFCONDUCT at multilingual counterspeech generation: A context-aware model for robust counterspeech generation in low-resource languages,” in Proceedings of the First Workshop on Multilingual Counterspeech Generation, 2025, pp. 37–46.
[13] M. Fanton, H. Bonaldi, S. S. Tekiro˘glu, and M. Guerini, ”Human-in-the-loop for data collection: a multi-target counter narrative dataset to fight online hate speech,” in Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), 2021, pp. 3226–3240.
[14] M. Trokhymovych and D. Saez-Trumper, ”Wikicheck: An end-to-end open source automatic factchecking api based on wikipedia,” in Proceedings of the 30th ACM International Conference on Information & Knowledge Management, 2021, pp. 4155–4164.
[15] Y. Wu, Z. Jin, C. Shi, P. Liang, and T. Zhan, ”Research on the application of deep learning-based BERT model in sentiment analysis,” arXiv preprint arXiv:2403.08217, 2024.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)