In proceedings of Workshop on Causal AI for Robust Decision Making (CARD 2025), held in conjunction with 23rd International Conference on Pervasive Computing and Communications (PerCom 2025)

Evaluating Causal AI Techniques for Health Misinformation Detection

Ommo Clark and Karuna P Joshi
Information Systems Department
University of Maryland Baltimore County, USA
{oclark1, karuna.joshi}@umbc.edu

***Abstract*—The proliferation of health misinformation on social media, particularly regarding chronic conditions such as diabetes, hypertension, and obesity, poses significant public health risks. This study evaluates the feasibility of leveraging Natural Language Processing (NLP) techniques for real-time misinformation detection and classification, focusing on Reddit discussions. Using logistic regression as a baseline model, supplemented by Latent Dirichlet Allocation (LDA) for topic modeling and K-Means clustering, we identify clusters prone to misinformation. While the model achieved a 73% accuracy rate, its recall for misinformation was limited to 12%, reflecting challenges such as class imbalance and linguistic nuances. The findings underscore the importance of advanced NLP models, such as transformer-based architectures like BERT, and propose the integration of causal reasoning to enhance the interpretability and robustness of AI systems for public health interventions.**

***Index Terms*—Health Misinformation, Misinformation Detection, Natural Language Processing, Social Media Analysis, Topic Modeling, Causal AI, Digital Health**

I. Introduction

Social media provides easy access to health information; however, it also facilitates rapid spread of misinformation [1]. Chronic conditions such as diabetes, hypertension, and obesity are particularly vulnerable to misleading health claims, which can erode trust in medical professionals and lead to adverse outcomes [2], [3]. Platforms like Reddit host a mix of credible advice and anecdotal misinformation, often cloaked in persuasive language [4], [5].

Traditional NLP methods, such as logistic regression, rely heavily on correlation-based patterns (for example, keywords), which may yield high accuracy but low recall for harmful misinformation. These black-box approaches lack interpretability, making it difficult to understand why certain posts are flagged or missed. Given the high stakes of health decisions, this study explores how causal reasoning can bolster NLP's ability to detect and classify misinformation more effectively [6].

We specifically investigate:

- Designing preprocessing workflows for text data.
- Training and evaluating supervised models for binary classification.
- Applying unsupervised techniques (topic modeling, clustering) to locate thematic and causal patterns.

By integrating causal insights, we aim to improve both the performance and interpretability of automated misinformation detection systems. In this paper, we discuss the related work in section II. Section III details our methodology. The results and discussion are given in Section IV, and we conclude in Section V.

II. Related Work

Health misinformation on social media has become a widespread concern, particularly after the COVID-19 infodemic [7]. Existing studies highlight the rapid spread of false or misleading health content on platforms such as Facebook, Twitter, and Reddit, where user-generated discussions can inadvertently propagate misinformed narratives [8]. Chronic conditions such as diabetes, hypertension, and obesity are especially prone to dubious claims and home remedies, due in part to their complexity and the human tendency to seek simple fixes [9]. This public health risk underscores the need for scalable misinformation detection systems capable of distinguishing credible from spurious content in real time.

Traditionally, misinformation detection has relied on correlation-based classification methods, e.g. logistic regression, support vector machines, or random forests that identify patterns in textual data by leveraging frequencies, n-grams, sentiment, or user features [10]. Although these approaches provide a baseline level of predictive performance, they often lack robust interpretability and do not capture the deeper cause-effect relationships that underlie the proliferation of misinformation. For example, correlation-based methods may note that certain keywords (for example 'natural', 'cure', 'remedy') co-occur more frequently in false claims, but do not clarify why these linguistic markers drive greater virality or user trust in misinformation [11].

In contrast, cause-effect-based detection aims to move beyond pattern recognition to model how misinformation arises, spreads, and influences user behavior. Bridging this gap is particularly important for real-time detection: If we understand why certain narratives go viral and which causal factors, such as emotional tone or social reinforcement, propel misinformation, then interventions (e.g. fact-check labels or post-removal labels) can be more precisely targeted and more effective [12]. Without a causal lens, systems risk overfitting to superficial textual cues and missing subtle yet potent misinformation that is context-dependent or evolves over time. Studies have shown that emotional language amplifies the virality of misinformation [13], making it more likely to be shared and believed. Conversely, interventions like fact-checking and prioritizing evidence-based content have been shown to causally reduce misinformation's influence on user behavior.

In recent years, Causal AI has gained traction as a way to improve not only the accuracy of prediction, but also the interpretability of machine learning models [14], [15]. Within Natural Language Processing (NLP), the integration of causal inference techniques typically involves two primary strategies: **S**tructural Causal Models (SCMs) in textual settings SCMs offer a formal way to represent and reason about the variables (nodes) and causal relationships (edges) that govern the dynamics of misinformation [**?**]. When applied to social media data, an SCM may include factors such as linguistic style, emotional appeal, or user engagement behaviors [16]. By structuring these factors into a directed acyclic graph (DAG), researchers can model the pathways through which misinformation influences user beliefs or spreads across communities.

- Operationalization: In a text-based setting, each post can be associated with potential causes (e.g., presence of pseudoscientific keywords, sensational language), mediators (e.g., user credibility, forum topic), and outcomes (e.g., number of shares, user trust). The SCM is learned or partially specified by domain experts, and interventions (e.g., removing misleading keywords or adding fact-check warnings) can be simulated to estimate their causal impact on user behavior.

Counterfactual Language Analysis Counterfactuals allow researchers to ask "what if?" questions by altering a specific aspect of the text or user interaction and measuring the resulting changes [17]. For instance, one can simulate scenarios where certain misinformation-heavy posts are removed or flagged to see if the user community shifts toward more evidence-based sources.

- Applications in Misinformation: By generating or modifying text to remove pseudoscientific claims (or inject factual sources), one can compare user reactions (engagement, sentiment) before and after the hypothetical edit. This approach not only reveals whether a post is likely to be misinformation, but also measures how an intervention might reduce its potential harm.

Alongside SCMs and counterfactual approaches, causal feature engineering has emerged as a complementary method. Rather than simply extracting correlation-based linguistic features (e.g. TF-IDF), researchers identify and encode variables that are hypothesized to have causal relevance for misinformation. Such features may include the author's domain expertise, the source's credibility, or specific terms with historically validated causal impacts (e.g., "magic pill," "miracle cure"). Incorporating these features can enhance recall for misinformation-laden posts by focusing on the textual and contextual signals that drive false narratives, rather than those that merely correlate with them [18].

Real-time detection of health misinformation necessitates both speed and accuracy, especially for public health interventions. Traditional correlation-based models can quickly classify large volumes of data but often stumble when misinformative content is linguistically subtle or novel. Causal models, while sometimes more complex to develop and update, can identify key drivers of misinformation and thus maintain higher recall and robustness over time. Additionally, causal explanations can guide moderators, healthcare professionals, and policy-makers in crafting targeted countermeasures—such as labeling suspicious posts or boosting credible content—thereby reducing the overall reach and impact of misinformation [16].

In summary, the integration of Structural Causal Models, counterfactual language analysis, and causal feature engineering represents a promising frontier for misinformation research. By explicitly modeling*why* and *how* misinformation propagates, these methods address the pitfalls of purely correlation-based approaches. This is particularly relevant in the health domain, where the cost of undetected misinformation can be measured in real-world consequences, such as delayed treatment or reduced trust in legitimate healthcare guidance.

III. Methodology

A. Overview of Overall Framework

We follow a systematic NLP pipeline, enhanced by causal inference techniques:

1) Data Preprocessing (cleaning, tokenization, lemmatization, TF-IDF)
2) Exploratory Data Analysis (distribution, engagement, topic mentions)
3) Model Development (baseline logistic regression, topic modeling, clustering)
4) Causal Approaches (counterfactual reasoning, causal feature engineering, heuristic-based removal)
5) Evaluation (accuracy, precision, recall, F1-score, and causal impact metrics)

B. Dataset

We extracted 437 posts from health-related subreddits (AskDocs, Diabetes, Health). Medical experts labeled each post as misinformation or accurate, leading to a 29:71 (misinformation vs. accurate) class imbalance. Key challenges included informal language, anecdotes, and varied post lengths.

C. Preprocessing

To prepare the data for modeling, we applied multiple steps to ensure consistency and remove noise, improve the quality of the text features and incorporate causal signals.

1) **Text Cleaning**: Removed stop words, punctuation, emojis, URLs.
2) **Tokenization & Lemmatization**: Converted words to base forms.
3) **Feature Engineering**:
    - TF-IDF vectors (5,000-term vocabulary).
    - *Causal Features*: Emotional tone flags, pseudoscientific keywords, user credibility signals.
4) **Class Imbalance Techniques**:

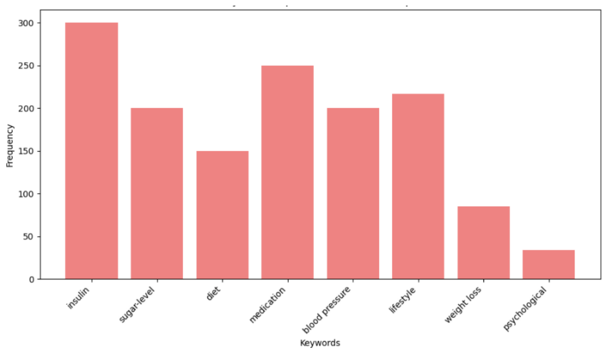

Fig. 1. Keyword frequencies within each topic

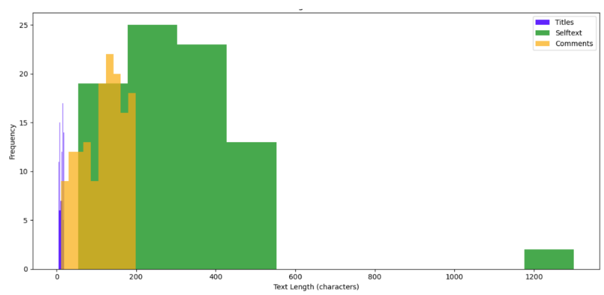

Fig. 2. distributions for titles, selftext, and comments

- SMOTE to oversample misinformation.
- Undersampling of the majority class.
- Weighted loss functions to penalize misclassification of misinformation.

D. Exploratory Data Analysis (EDA)

EDA was a crucial step in understanding the structure, distribution, and nuances of the dataset [19]. This study uses a dataset of Reddit posts related to diabetes, hypertension, and obesity, annotated as ”misinformation” or ”accurate information” by medical experts to explore the patterns and themes in the dataset and provide a foundation for the model development and evaluation.

Topics of interest included diabetes (953 mentions), hypertension (667 mentions), and obesity (119 mentions). Posts varied in length; titles were generally short, while longer selftext entries often interwove accurate and misleading claims. This complexity underscored the importance of advanced feature engineering to capture subtle cues.

IV. Experimental Results and Evaluation

The results highlight the performance of the machine learning models, thematic patterns from unsupervised learning, and key challenges and opportunities for improvement.

A. Logistic Regression Results

The baseline logistic regression model provided insights into the dataset’s challenges and model limitations:

1) Accuracy: The model correctly classified 73% of all posts, but this metric alone is insufficient due to class imbalance.

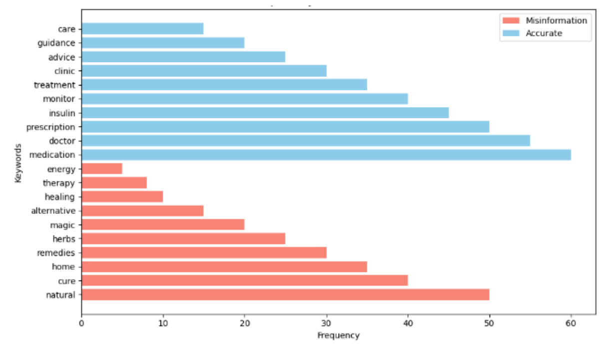

Fig. 3. Frequently used keywords

Metrics	Value
Accuracy	73%
Precision (Misinformation)	75%
Recall (Misinformation)	12%
F1-Score (Misinformation)	20%

TABLE I
THEMES AND DESCRIPTIONS

2) Precision for Misinformation: Of all posts predicted as misinformation, 75% were correct. This indicates that the model minimizes false positives for misinformation.
3) F1-Score for Misinformation: At 20%, this metric highlights the imbalance between precision and recall for misinformation detection.

Confusion Matrix:

Actual vs. Predicted	Misinformation	Accurate
Misinformation	3 (True Positives)	23 (False Negatives)
Accurate	1 (False Positives)	61 (True Negatives)

TABLE II
CONFUSION MATRIX

B. Topic Modeling (LDA)

Latent Dirichlet Allocation (LDA) was used to extract themes from the dataset [20], revealing underlying patterns

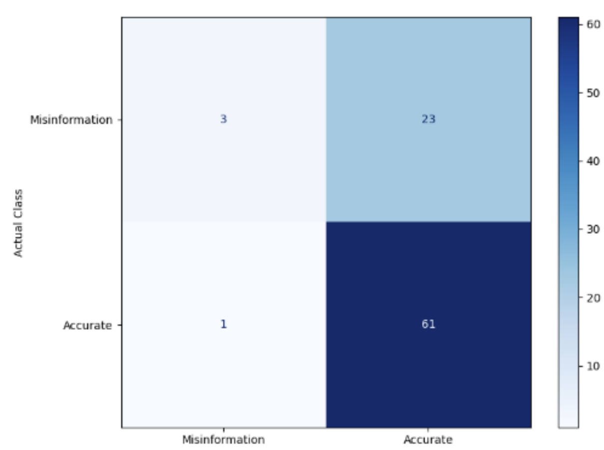

Fig. 4. Confusion Matrix

Fig. 5. HeatMap

in the discussions.

Key Topics Identified:

1) Diabetes Management:
   a) Keywords: ”insulin,” ”diet,” ”blood sugar.”
   b) Discussions centered around managing blood sugar levels, insulin use, and dietary advice.

b) Hypertension Control:
   a) Keywords: ”blood pressure,” ”medication,” ”doctor.”
   b) Topics included medication regimens, self-monitoring tools, and consultations.

c) Home Remedies and Alternative Treatments:
   a) Keywords: ”natural,” ”herbs,” ”remedy.”
   b) Posts highlighted anecdotal advice, often containing misinformation.

d) Health Monitoring:
   a) Keywords: ”wearables,” ”tracking,” ”device.”
   b) Focus on self-monitoring tools like blood pressure cuffs and glucose monitors.

e) Emotional Support and Personal Stories:
   a) Keywords: ”experience,” ”support,” ”coping.”
   b) Posts shared emotional journeys and coping mechanisms.

Coherence Score: The LDA model achieved a coherence score of 0.21, indicating topics are moderately coherent but could be improved with better preprocessing and topic refinement. **I**nsights:

a) Misinformation is concentrated in topics like “home remedies,” reflecting the need for targeted misinformation detection in these areas.
b) Emotional support posts, while generally accurate, may propagate misinformation unintentionally.

C. Clustering Results (K-Means)

K-Means clustering grouped posts into thematic clusters based on their textual similarity.

Insights:

Cluster	Description
Cluster 0	Posts discussing alternative remedies, anecdotal advice and unverified health advice (misinformation-heavy)
Cluster 1	Evidence-based discussions, such as doctor consultations and medication usage, and proven treatments (accuracy-heavy).
Cluster 2	Posts sharing personal stories, experiences and emotional support.
Cluster 3	Discussions on health monitoring tools, wearable devices and self-measured metrics.
Cluster 4	General health discussions about lifestyle and diet.

TABLE III
CLUSTERS IDENTIFIED

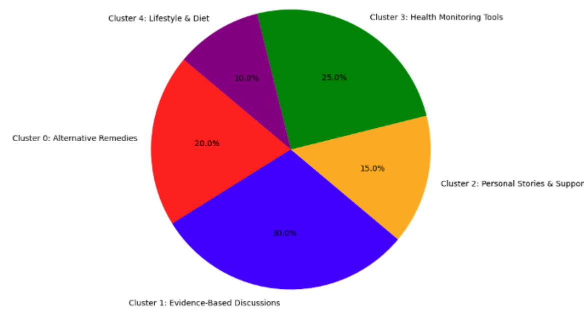

Fig. 6. Evidence Based Discussions

a) Certain clusters (e.g., Cluster 0) are more likely to contain misinformation, suggesting targeted public health interventions
b) Improved clustering methods (e.g., DBSCAN) or better feature representations (e.g., BERT embeddings) could enhance cluster separations
c) Third item

D. Counterfactual Analysis & Heuristics

A simple heuristic-based approach flagged posts with high predicted probabilities of misinformation. Removing these “high-risk” posts in a simulated environment boosted engagement with accurate content. This under-

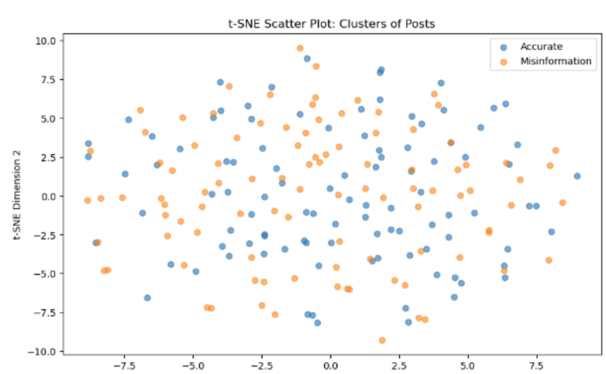

Fig. 7. Clusters of Posts

scores the potential for more formal SCMs or counterfactual text generation to guide real-time interventions.

E. Engagement Patterns

a) Text Length Analysis: Posts with longer selftext are more likely to include nuanced discussions, often blending accurate information with misinformation. Titles are concise, with an average length of 10–12 words, summarizing the main themes.

b) Insights: Longer posts tend to attract more engagement (comments), suggesting higher exposure and potential for misinformation spread.

Discussion

This study demonstrates that while correlation-based methods like logistic regression achieve moderate accuracy (73%), they often fail to detect nuanced health misinformation, as evidenced by the low recall of 12%. Topic modeling and clustering uncovered specific misinformation "hotspots," such as posts promoting "natural cures" and pseudoscientific claims, underscoring the need for more advanced classification strategies.

Health misinformation on social media is a public health crisis with tangible consequences [21]. Community-driven platforms like Reddit host two predominant content types: fact-based posts grounded in medical evidence and anecdotal misinformation rooted in personal experiences or pseudoscientific beliefs [22]–[24]. Claims about herbal teas curing diabetes, for example, are compelling because they promise simple solutions to complex problems, leveraging emotional appeal and the veneer of "scientific" language. This relatability and persuasive tone make misinformation difficult to counter with traditional NLP approaches.

Class imbalance was a core limitation, even after oversampling via SMOTE. When misinformative posts comprise only 29% of the dataset, subtle cues—like "natural cures" or unreferenced claims of scientific backing—become harder to distinguish from legitimate medical advice. Overlapping linguistic patterns in posts discussing diets or alternative remedies added further complexity, highlighting the importance of incorporating context- and meaning-focused approaches, as opposed to solely relying on surface-level keyword matching [25], [26].

The findings also show how emotionally charged narratives can amplify misinformation's persuasiveness. Posts using terms like "struggle," "hope," or personal success stories can inadvertently frame anecdotal evidence as universal truths. From a public health perspective, every overlooked misinformation post represents a risk, either leading to poor patient decisions or shifting community perceptions about chronic disease treatments.

Logistic regression's shortfalls suggest a need for more sophisticated models capable of capturing deeper linguistic and contextual nuances [27]. Transformer-based architectures such as BERT show promise in capturing contextual relationships more effectively [28]. Future iterations could leverage these models to improve recall, differentiate factual from anecdotal narratives, and incorporate metadata (e.g., user engagement) to provide a holistic view of each post's credibility.

Moreover, embedding causal signals—like emotional tone, user credibility, or pseudoscientific markers—into NLP pipelines provides valuable interpretability. In practical, human-AI collaboration, such causal insights allow healthcare professionals to see why certain posts are flagged and to offer evidence-based counterpoints. This fosters trust in AI-driven systems, as moderators and medical experts can directly address the emotional or linguistic cues fueling misinformation.

Ultimately, this research underscores the urgent need for interdisciplinary approaches, bringing together AI practitioners, healthcare professionals, and social scientists. As social media shape public perceptions of health, the tools developed here lay the groundwork to mitigate the spread of misinformation. They highlight how seemingly anecdotal or harmless content can carry far-reaching implications for public health.

Conclusion

Our findings confirm the feasibility of integrating causal reasoning with NLP to detect health misinformation more reliably. By focusing AI decisions on cause-effect dynamics, this framework can improve interpretability, reduce missed misinformation cases, and support real-time interventions. Future work will focus on deploying real-time causal inference systems that dynamically analyze misinformation spread, employing SCMs and counterfactual reasoning to improve interventions. Furthermore, these causal frameworks will be extended to other high-stakes domains such as disaster response and financial risk assessments, aiming to strengthen collaboration between AI systems and human decision makers in critical decision-making processes.

Future Work

This study demonstrates the feasibility of using causal AI for misinformation detection and lays a solid foundation for future advancements. As we continue to refine and expand this work, several promising directions emerge.

One key area for improvement is broadening the scope within pervasive computing. Although this research provides valuable information on misinformation detection, future work will explore how device-level interactions and communication-layer dynamics can further improve misinformation detection and intervention strategies in ubiquitous real-time environments.

We will also improve the recall performance for misinformation detection. The current model effectively identifies misinformation but faces challenges due to dataset imbalances. To address this, we will implement advanced class balancing techniques to improve recall while maintaining high precision.

Furthermore, while this study successfully integrates causal AI into misinformation detection, there is an opportunity to further innovate within causal reasoning methodologies. Future work will focus on developing domain-specific Structural Causal Models (SCMs) and refining counterfactual analysis techniques, ensuring that misinformation detection benefits from deeper causal insights and improved interpretability.

Moreover, we will enhance the scalability and robustness of misinformation detection by integrating transformer-based architectures (e.g., DistilBERT) and knowledge graph-based causal inference. These advances will enable more context-sensitive analysis, leading to a more accurate classification of subtle and evolving misinformation.

Building on these advancements, this research will expand its impact beyond NLP-based misinformation detection, contributing to a comprehensive, causally aware framework for misinformation analysis in pervasive computing and beyond.

Acknowledgment

We acknowledge the assistance of medical experts in annotating the project dataset.

References

[1] O. Clark, K. P. Joshi, T. Reynolds *et al.*, "Global relevance of online health information sources: A case study of experiences and perceptions of nigerians," in *AMIA 2024 Annual Symposium*, 2024.

[2] T. Burki, "Social media and misinformation in diabetes and obesity," *The Lancet Diabetes & Endocrinology*, vol. 10, no. 12, p. 845, 2022.

[3] R. Patil, S. S. Datta, A. V. Boratne *et al.*, "Popular misconceptions regarding the diabetes management: where should we focus our attention?" *Journal of Clinical and Diagnostic Research: JCDR*, vol. 7, no. 2, p. 287, 2013.

[4] J. Chen, Y. Wang *et al.*, "Social media use for health purposes: systematic review," *Journal of medical Internet research*, vol. 23, no. 5, p. e17917, 2021.

[5] L. Bozarth, J. Im, C. Quarles, and C. Budak, "Wisdom of two crowds: Misinformation moderation on reddit and how to improve this process—a case study of covid-19," *Proceedings of the ACM on Human-Computer Interaction*, vol. 7, no. CSCW1, pp. 1–33, 2023.

[6] A. Feder, K. A. Keith, E. Manzoor, R. Pryzant, D. Sridhar, Z. Wood-Doughty, J. Eisenstein, J. Grimmer, R. Reichart, M. E. Roberts *et al.*, "Causal inference in natural language processing: Estimation, prediction, interpretation and beyond," *Transactions of the Association for Computational Linguistics*, vol. 10, pp. 1138–1158, 2022.

[7] D. N. I. BORGES, I. Nascimento, A. Pizarro, J. Almeida, N. AZZOPARDI-MUSCAT, M. GONÇALVES, M. BJÖRKLUND, and D. NOVILLO-ORTIZ, "Infodemics and health misinformation: a systematic review of reviews," *Bulletin of the World Health Organization*, vol. 100, no. 9, pp. 544–561, 2022.

[8] S. Kisa and A. Kisa, "A comprehensive analysis of covid-19 misinformation, public health impacts, and communication strategies: Scoping review," *Journal of Medical Internet Research*, vol. 26, p. e56931, 2024.

[9] O. Clark, T. L. Reynolds, E. C. Ugwuabonyi, and K. P. Joshi, "Exploring the impact of increased health information accessibility in cyberspace on trust and self-care practices," in *Proceedings of the 2024 ACM Workshop on Secure and Trustworthy Cyber-Physical Systems*, 2024, pp. 61–70.

[10] V. S. Pendyala, "Misinformation containment using nlp and machine learning: Why the problem is still unsolved," in *Deep Learning Research Applications for Natural Language Processing*. IGI Global, 2023, pp. 41–56.

[11] Z. Zhou, H. Guan, M. M. Bhat, and J. Hsu, "Fake news detection via nlp is vulnerable to adversarial attacks," *arXiv preprint arXiv:1901.09657*, 2019.

[12] Z. Chen, L. Hu, W. Li, Y. Shao, and L. Nie, "Causal intervention and counterfactual reasoning for multi-modal fake news detection," in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2023, pp. 627–638.

[13] K. Solovev and N. Pröllochs, "Moral emotions shape the virality of covid-19 misinformation on social media," in *Proceedings of the ACM web conference 2022*, 2022, pp. 3706–3717.

[14] M. O. Gani, S. Kethireddy, R. Adib, U. Hasan, P. Griffin, and M. Adibuzzaman, "Structural causal model with expert augmented knowledge to estimate the effect of oxygen therapy on mortality in the icu," *Artificial intelligence in medicine*, vol. 137, p. 102493, 2023.

[15] M. Nirav Shah and A. Ganatra, "A systematic literature review and existing challenges toward fake news detection models," *Social Network Analysis and Mining*, vol. 12, no. 1, p. 168, 2022.

[16] M.-Y. Chen, Y.-W. Lai, and J.-W. Lian, "Using deep learning models to detect fake news about covid-19," *ACM Transactions on Internet Technology*, vol. 23, no. 2, pp. 1–23, 2023.

[17] Z. Chen, Q. Zhang, J. Liu, Y. Wang, H. Lv, L. Wang, J. Duan, M. Xv, and H. Wang, "Counterfactual multimodal fact-checking method based on causal intervention," in *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*. Springer, 2024, pp. 582–595.

[18] W. Talukdar and A. Biswas, "Synergizing unsupervised and supervised learning: A hybrid approach for accurate natural language task modeling," *arXiv preprint arXiv:2406.01096*, 2024.

[19] S. Vali Zade and H. Rastegar, "Parallel coordinates plot: A visual examination of data structures in exploratory data analysis," *Analytical and Bioanalytical Chemistry Research*, vol. 11, no. 2, pp. 201–209, 2023.

[20] S. Zhou, P. Kan, Q. Huang, and J. Silbernagel, "A guided latent dirichlet allocation approach to investigate real-time latent topics of twitter data during hurricane laura," *Journal of Information Science*, vol. 49, no. 2, pp. 465–479, 2023.

[21] M. Shahbazi and D. Bunker, "Social media trust: Fighting misinformation in the time of crisis," *International Journal of Information Management*, vol. 77, p. 102780, 2024.

[22] A. Grasser-Parger, "Community-driven design: Collaborative objects, community-generated content, and the expanded role of architects," *International Journal of Architectural Computing*, p. 14780771241234450, 2024.

[23] L. M. Malki, D. Patel, and A. Singh, "" the headline was so wild that i had to check": An exploration of women's encounters with health misinformation on social media," *Proceedings of the ACM on Human-Computer Interaction*, vol. 8, no. CSCW1, pp. 1–26, 2024.

[24] F. Pilati, T. Venturini, P. L. Sacco, and F. Gargiulo, "Pseudo-scientific versus anti-scientific online conspiracism: A comparison of the flat earth society's internet forum and reddit," *New Media & Society*, p. 14614448241252593, 2024.

[25] T. Chani, O. O. Olugbara, and B. Mutanga, "The problem of data extraction in social media: A theoretical framework," *Journal of Information Systems and Informatics*, vol. 5, no. 4, pp. 1363–1384, 2023.

[26] S. Van der Linden, *Foolproof: Why misinformation infects our minds and how to build immunity*. WW Norton & Company, 2023.

[27] Á. I. Rodríguez and L. L. Iglesias, "Fake news detection using deep learning," *arXiv preprint arXiv:1910.03496*, 2019.

[28] A. Onan, "Hierarchical graph-based text classification framework with contextual node embedding and bert-based dynamic fusion," *Journal of king saud university-computer and information sciences*, vol. 35, no. 7, p. 101610, 2023.