British Journal of Multidisciplinary and Advanced Studies,6(3)1-31, 2025
Engineering and Technology
Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

SENS-HEAD: A Machine Learning Framework for Sensationalism Detection in News Headlines Using Linguistic and Semantic Features

Po-Hsuan Chang [1], Akshi Kumar[1]*, Saurabh Raj Sangwan[2]
[1]School of Computing, Goldsmiths, University of London, London, United Kingdom
[2] Department of Computer Science and Engineering, Artificial Intelligence and Machine Learning, G L Bajaj Institute of Technology and Management, Greater Noida, India
*Akshi.Kumar@gold.ac.uk

doi: https://doi.org/10.37745/bjmas.2022.04909 Published June 01, 2025

Citation: Chang P., Kumar A., Sangwan SR (2025) SENS-HEAD: A Machine Learning Framework for Sensationalism Detection in News Headlines Using Linguistic and Semantic Features, *British Journal of Multidisciplinary and Advanced Studies*,6(3)1-31

Abstract: *The proliferation of sensationalized news headlines has raised concerns about media integrity, necessitating automated approaches for detecting sensationalism beyond traditional clickbait classification. This study presents SENS-HEAD, a novel dataset comprising over 30,000 annotated headlines labelled for sensational content and emotional arousal. Employing Natural Language Processing (NLP), we extract a diverse set of linguistic and semantic features, including sentiment polarity, syntactic complexity, punctuation distribution, and stop word ratio, to systematically distinguish sensational from non-sensational headlines. We implement ensemble learning models—XGBoost, CATBoost, and Random Forest achieving a balanced F1-score of 0.66. To enhance interpretability, we integrate SHAP (SHapley Additive exPlanations), unveiling key predictive markers such as stop word frequency, headline length, and sentiment extremity. The findings not only advance explainable AI (XAI) for sensationalism detection but also provide practical applications in automated journalism, content moderation, and media ethics regulation. By strengthening computational linguistics with ethical AI, this research delivers actionable insights for policymakers and promotes trustworthy news dissemination in the digital era.*

Keywords: Sensationalism detection, linguistic features, NLP, machine learning, news headlines, XAI

INTRODUCTION

In the digital age, news headlines play a critical role in capturing readers' attention and driving engagement. The transition from traditional print media to online platforms has intensified competition among news outlets, incentivizing the use of sensationalism to attract more clicks and boost traffic. Sensationalism, characterized by emotionally charged, exaggerated, or dramatic language, is designed to provoke strong reactions from readers. While sensational headlines are effective in increasing click-through rates (CTR), they often distort the reality of events by prioritizing engagement over accuracy [1]. This practice raises ethical concerns about the dissemination of information and its long-term impact on public perception. Despite these concerns, the detection of sensationalism in news headlines has been underexplored, with most research focusing on clickbait detection.

To address this gap, this paper presents an automated system for detecting sensationalism in news headlines using a linguistic feature-based approach. The foundation of this study is the SENS-HEAD dataset, a novel resource developed specifically for the task of sensationalism detection. The dataset was recreated from existing clickbait data[1], meticulously annotated to focus on sensationalism rather than simply misleading content. Comprising over 30,000 annotated headlines, SENS-HEAD serves as the basis for training machine learning models to classify sensationalism in headlines based on linguistic patterns. The motivation for this research is driven by ethical concerns surrounding the rise of sensationalism in media and its influence on public perception. While clickbait has been widely studied [2, 3], sensationalism represents a broader issue that extends beyond simple click-maximization strategies. Sensational headlines often manipulate readers' emotions by appealing to curiosity, fear, or excitement, shaping the way audiences perceive events in ways that may not reflect the true content. Unlike clickbait, which focuses primarily on luring readers into clicking through misleading cues, sensationalism engages deeper emotional responses by using specific linguistic strategies [4, 5]. Therefore, detecting sensationalism is crucial for ensuring the accuracy and integrity of news reporting.

Consider the example of a headline like *"Shocking! Man Survives Lightning Strike for the Third Time."* The use of the word "shocking" and the exclamation mark conveys a sense of drama and surprise, which heightens emotional engagement. In contrast, a non-sensational version might simply state, *"Man Survives Third Lightning Strike."* Both headlines present the same factual information, but the sensational version is crafted to provoke a stronger emotional response, thus drawing more attention. Similarly, a headline like *"World's Best Athlete Breaks Record!"* uses a superlative and exclamation mark to exaggerate excitement, while a more factual version, *"Athlete Breaks Record in Championship,"* provides the information without evoking unnecessary drama. These examples highlight how subtle linguistic choices, such as exaggerated adjectives and punctuation, can significantly influence reader engagement and emotional reaction.

This research adopts a natural language processing (NLP) and machine learning framework to analyze and detect sensationalism based on these linguistic patterns. Various features, such as sentiment analysis, readability scores, punctuation usage, and syntactic complexity, are extracted from headlines to form the input for machine learning models. The models used in this study—XGBoost, CATBoost, and Random Forest—are known for their robust performance in classification tasks. To further enhance interpretability, SHAP (SHapley Additive exPlanations) [6] analysis is applied to understand the contribution of each linguistic feature to the model's predictions, identifying the most influential markers of sensationalism. The primary contribution of this research is threefold:

- **Creation of the SENS-HEAD dataset**: A comprehensive dataset comprising 30,000+ annotated headlines, explicitly curated for sensationalism detection, surpassing conventional clickbait classification by incorporating emotional arousal and linguistic markers
- **Development of an NLP-Driven Machine Learning Framework:** A robust classification system leveraging linguistic and semantic features including sentiment polarity, punctuation patterns, syntactic complexity, and readability metrics to accurately differentiate sensational and non-sensational headlines.
- **Integration of explainability through SHAP (SHapley Additive exPlanations) analysis**: Enhancing model transparency, SHAP (SHapley Additive exPlanations) identifies key linguistic indicators contributing to sensationalism classification. This interpretability makes

[1] https://www.kaggle.com/datasets/amananandrai/clickbait-dataset

the approach suitable for practical deployment in content moderation, media ethics enforcement, and automated journalism, supporting responsible AI-driven news evaluation.

The rationale for adopting a linguistic feature-based approach is grounded in the observation that language is the primary medium through which sensationalism is communicated. By focusing on the structural and emotional components of language, we can identify patterns that signal sensationalism more effectively than superficial metrics such as headline length or engagement statistics. Additionally, existing research on clickbait detection, while related, does not fully account for the emotional and psychological manipulation inherent in sensationalism. Clickbait detection primarily examines intent whether the headline misleads the reader into clicking whereas sensationalism goes beyond intent, targeting emotional engagement and altering how readers perceive the world. Addressing these complexities not only fills a gap in current literature but also offers practical applications in content moderation and media regulation.

By utilizing advanced machine learning models and SHAP analysis, this study offers a transparent and interpretable framework for detecting sensationalism. SHAP allows us to understand how individual linguistic features, such as sentiment polarity or punctuation patterns, influence the model's predictions. This interpretability is crucial for applying the findings in real-world settings, where ethical considerations must balance attention-grabbing headlines with responsible reporting.

In summary, this research provides a comprehensive approach to sensationalism detection, offering valuable insights for journalists, editors, and regulators. It not only builds upon existing knowledge of headline detection but also paves the way for future applications in automated journalism and media oversight. The insights gained from this study can help create more balanced and ethical media practices, ensuring that attention-grabbing content does not come at the expense of truthfulness and public trust.

The organization of the paper follows a logical progression, starting with an introduction that establishes the importance of detecting sensationalism in headlines and the rationale for focusing on linguistic features. This is followed by a literature review that provides context from past studies, positioning the paper within the existing body of research. The methodology section describes the development of the SENS-HEAD dataset, feature extraction, and the multi-stage model training process. A detailed discussion on the ensemble machine learning models and SHAP analysis is presented, highlighting model explainability. Finally, the paper concludes with a discussion of results, limitations, and directions for future work.

LITERATURE REVIEW

Sensationalism in headlines has long attracted interest from both linguistic and psychological researchers due to its significant influence on public perception and emotional engagement. Over time, studies have evolved to explore how sensationalism is communicated through language, employing various methods from rhetorical analysis to machine learning-based approaches. This section reviews the key literature on sensationalism, tracing its development from early studies to the most recent advances in 2023–2024. The review highlights existing gaps and lays the foundation for the current research.

Early research, such as that by Molek-Kozakowska [4], examined how sensational headlines use rhetorical strategies, including hyperbole, exaggeration, and dramatic language, to elicit strong cognitive and emotional responses from readers. Molek-Kozakowska analyzed the linguistic devices that amplify a story’s emotional resonance, such as the use of extreme adjectives or verbs, demonstrating how these features shape readers' perceptions of events. Similarly, Kleemans and Vettehen [7] linked sensationalism to evolutionary theory, arguing that the human brain has adapted to scanning for information related to survival and reproduction. This explains the long-standing appeal of sensationalism in news stories, as such content resonates with instinctual human responses to threats and opportunities.

The methodological focus on sensationalism detection shifted with the growing popularity of Natural Language Processing (NLP) techniques and the application of machine learning in text analysis. Biyani et al. [8], in 2016, introduced a novel approach by utilizing linguistic markers, such as superlative adjectives and specific punctuation (e.g., exclamation marks), in their clickbait detection models. Their study highlighted the reliability of certain linguistic features in identifying sensational content. Brown et al. [9] explored sensationalism as a broader concept, often tied to emotional appeal and shock value, and examined how it affects audience engagement on social media across various countries. Additionally, Montejo and Adriano [10] analyzed how Philippine news headlines use linguistic strategies, such as evaluative language and references to elite figures, to sensationalize stories and reflect media ideologies. Their use of Fairclough's Critical Discourse Analysis (CDA) framework revealed that sensationalism in headlines is frequently achieved through emotionally charged language, superlatives, and references to well-known figures to capture attention.

In 2019, Xu et al. [11] introduced a novel approach for sensational headline generation using reinforcement learning, proposing an Auto-tuned Reinforcement Learning (ARL) mechanism that dynamically balances maximum likelihood estimation with reinforcement learning. This innovation optimized both sensationalism and fluency in generated headlines without requiring labeled data. Mourão and Robertson [12] examined fake news as a form of discursive integration, analyzing the genre-blending of traditional journalism with misinformation, sensationalism, and clickbait across 50 sites. They found that moderate sensationalism and partisanship, rather than outright fabrications, drive engagement on social media. Arbaoui and Van der Brug [13], in 2020, provided a comparative analysis of sensationalism in news coverage across 14 television systems, revealing significant variations in sensationalist reporting patterns across different media contexts. In the same year, Naeem et al. [14] presented a deep learning framework for clickbait detection on social networks, combining a Part of Speech Analysis Module (POSAM) with an LSTM classifier, achieving a classification accuracy of 97%.

The COVID-19 pandemic saw a significant increase in sensationalism in media coverage across various regions. In 2021, Allen and Ayalon [15] critically analyzed how residential care during the COVID-19 pandemic was portrayed in leading American newspapers. They uncovered how sensationalist language and the exclusion of residents' voices perpetuated narratives of fear, danger, and distrust toward long-term care facilities. Similarly, Wasserman et al. [16] analyzed South African newspaper coverage of COVID-19, finding that alarmist and sensationalist reporting predominated, with nearly half of the front-page reports using alarmist narratives. This sensationalist framing amplified public anxieties and often failed to provide actionable health information, ultimately exacerbating public fears. Ottwell et al. [17] emphasized the prevalence of sensational media reporting about COVID-19 therapies, detection methods, and vaccines, often exaggerating the potential benefits without supporting clinical data.

In the same year, Liepukhova and Shcherbak [18] examined lexical and semantic strategies used in German-language media headlines to create sensationalism. They highlighted how expressive vocabulary and stylistic devices such as metaphors, allusions, and antithesis were employed to evoke strong emotional responses. In 2023, Esparza et al. [19] demonstrated the effectiveness of BERT-based classifiers in detecting sensational health-related headlines in Spanish, achieving up to 94% accuracy and F1-Score on nearly 2,000 Mexican newspaper headlines. More recently, Hamby et al. [20], in 2024,

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

investigated narrative differences between real and fake news, finding that fake news often features higher sentiment volatility and sensational storytelling, which significantly increases its spread on social media. Khawar and Boukes [21] analyzed sensationalism in news promotion on Twitter (now X), comparing legacy news outlets (e.g., USA Today) with online-native outlets (e.g., BuzzFeed). Their study revealed that online-native outlets employed more sensationalist features, positively correlating with higher user engagement, particularly in the form of likes and retweets.

Despite significant advances in sensationalism and clickbait detection, much of the research has focused primarily on clickbait, with limited efforts dedicated to sensationalism detection specifically. Furthermore, while there have been quantitative studies on media sources concerning sensational news, no comprehensive language or discourse analysis has been conducted to systematically examine sensationalism as a distinct phenomenon. Studies by Biyani et al. [8] and Brown et al. [9] enhanced understanding of the linguistic markers of sensationalism, yet these efforts lacked a robust framework that distinguishes sensationalism from other forms of emotional manipulation, such as clickbait. With the rise of deep learning and explainable AI models, the need for a focused dataset and methodologies addressing sensationalism in a more detailed and structured way has become evident. The paper by Zhang and Kejriwal [22], offers critical insights into the challenges of sensationalism detection in dynamic news environments, particularly highlighting the phenomenon of concept drift. Our work builds on these insights, offering a more comprehensive and interpretable framework with broader domain coverage and improved generalizability.

METHODOLOGY

To accurately detect and analyze sensationalism in news headlines, a methodologically sound and data-driven approach is critical. The methodology employed in this research is designed to systematically capture the nuances of sensationalism by leveraging linguistic features that trigger emotional arousal in readers. This approach not only ensures that the classification models are robust but also provides explainability, allowing us to understand the underlying factors contributing to sensationalism. The decision to adopt this feature-based NLP approach stems from the need to go beyond surface-level text analysis and delve into the structural and emotional components of language, which are essential for identifying sensational content. Figure 1 illustrates the workflow of the methodology adopted.

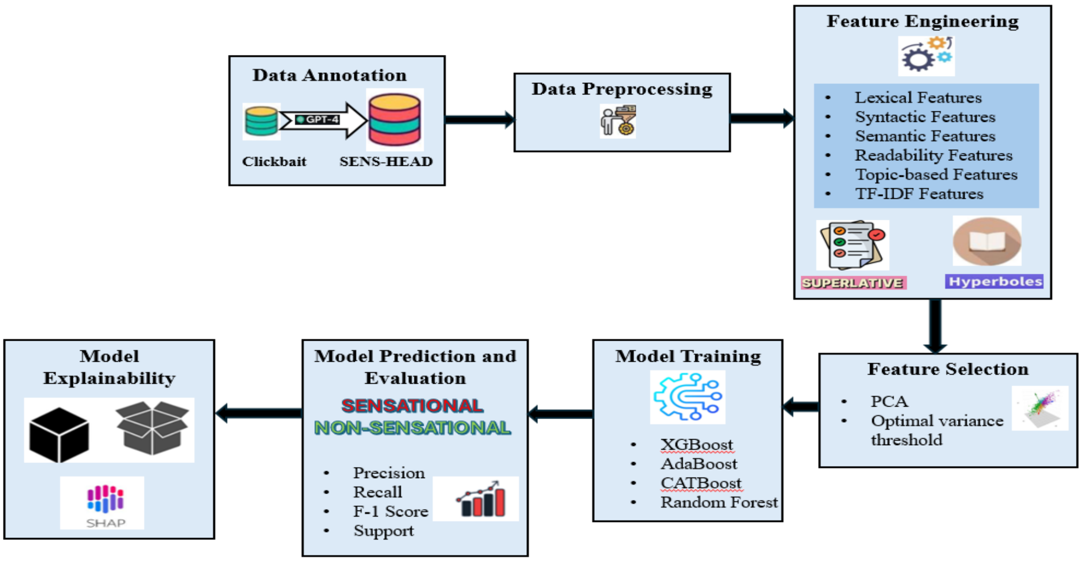

Fig.1. Overview of the Methodology for Detecting Sensationalism in News Headlines

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

The SENS-HEAD Dataset

The SENS-HEAD dataset[2] was developed to address the growing need for resources tailored specifically to detect sensationalism in news headlines. It combines diverse sources of clickbait and non-clickbait headlines, enabling a nuanced analysis of sensationalism across journalistic contexts. The dataset incorporates headlines from clickbait platforms (e.g., BuzzFeed, Upworthy, ViralNova, ThatScoop, ScoopWhoop, and ViralStories) and traditional outlets (e.g., WikiNews, The New York Times, The Guardian, and The Hindu). This diversity represents the two extremes of sensationalism, the exaggerated, emotion-driven content of clickbait versus the factual, balanced reporting of traditional journalism. By capturing these variations, the dataset allows for the development of a comprehensive classification model that reflects the entire spectrum of sensationalism tactics.

The SENS-HEAD dataset consists of 30,424 rows, reduced from the original 32,000 after removing 117 rows with conflicting sensation indicators and 1,459 rows with inconsistent annotations. This size ensures robustness by providing sufficient examples for training machine learning models, improving their ability to generalize to unseen data and achieve reliable statistical significance. The mix between clickbait and non-clickbait examples was carefully balanced to ensure fair representation of both classes, reducing bias in training and enhancing the model's ability to distinguish sensationalism effectively.

Creating the SENS-HEAD dataset involved several challenges:

- *Conflicting Annotations:* Resolving inconsistencies in annotations required careful analysis and refinement of scoring criteria.
- *Subjective Interpretation of Sensationalism:* Sensationalism often depends on context and perception, necessitating clear guidelines to ensure consistent labelling.
- *Balancing Strategies:* Maintaining a balanced dataset required ensuring equal representation of both sensational and neutral headlines, even as outliers were removed to improve quality.

Despite these challenges, the resulting dataset provides a high-quality resource for developing reliable and interpretable models for sensationalism detection.

Annotation Process

To ensure high-quality annotations, the dataset was processed using OpenAI's GPT-4 model, configured for advanced text analysis. The annotation task involved scoring headlines for sensationalism and emotional arousal based on a 5-point Likert scale [23], where 0 indicated "Not at all sensational" and 4 indicated "Very sensational." The model was guided using a few-shot prompting technique, which involved providing annotated examples to inform subsequent predictions. Human oversight was employed to review and validate the annotations, ensuring consistency and minimizing errors. Each headline was scored according to the following Likert scale:

(0): Not at all sensational (0-0.75)
(1): Not too much (0.76-1.50)
(2): Somewhat sensational (1.51-2.25)
(3): Fairly sensational (2.26-3.25)
(4): Very sensational (3.26-4)

The process was automated and conducted in batches of 10 headlines with a 5-second delay between batches to manage API limitations and maintain uniformity. This automation significantly enhanced efficiency, enabling the consistent evaluation of all 30,424 headlines for both sensationalism and emotional arousal. During the annotation process, special care was taken to handle ambiguous headlines and conflicting annotations. A set of rules was developed to ensure that the model's predictions aligned

[2] SENS-HEAD.xlsx

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

with linguistic markers and contextual cues indicative of sensationalism, such as hyperbolic language, exclamatory punctuation, and emotionally charged phrases.

Data Fields

The SENS-HEAD dataset consists of 10 detailed columns, each representing a specific attribute associated with the headlines. These features were extracted and annotated through a systematic process, ensuring that each headline was accurately assessed for its sensationalism and emotional impact. Table 1 summarizes the key dataset fields, their descriptions, and the processes used for annotation in the dataset:

Table 1. SENS-HEAD Dataset Fields and Annotation Processes

Field	Description	Process Used
Headline	The raw news headline, extracted from various sources in the original Kaggle dataset, including clickbait (e.g., *BuzzFeed*) and non-clickbait platforms (e.g., *The Guardian*).	Sourced from Kaggle's original clickbait dataset.
Clickbait Status	Classified each headline as either clickbait or non-clickbait based on its source. Clickbait platforms were labelled as clickbait, while traditional sources were labelled as non-clickbait.	Sourced from Kaggle's original clickbait dataset.
Sensation Score	Assigned a score from 0 to 4 for sensationalism, determined using linguistic markers like hyperbole, superlatives, and emotionally charged phrases.	Annotated using OpenAI's GPT-4, reviewed by human annotators.
Sensation Reason	Provided the rationale behind the sensation score based on elements such as exaggerated claims, emotional triggers (e.g., fear, excitement), and rhetorical devices.	GPT-4 annotation with human oversight.
Emotion	The primary emotion evoked, categorized into fear, anger, shock, excitement, or curiosity. Emotion detection models and GPT-4 were used for final classification.	GPT-4 assisted, human review for accuracy.
Arousal	Represents the emotional intensity, from low to high, triggered by the headline, based on linguistic features like exclamations, urgency, and charged adjectives.	Annotated using OpenAI's GPT-4, reviewed by human annotators.
Arousal Score	Numeric score (0 to 4) indicating arousal intensity. Calculated using text analysis algorithms that assess sentiment strength and emotional impact.	Calculated using NLP models and GPT-4, manually verified.
Arousal Reason	Justifies the arousal score based on language intensity, word choice, and rhetorical devices that increase emotional impact.	Manually annotated with human verification, supported by GPT-4.
Arousal Category	Categorized the type of arousal (e.g., curiosity, fear, excitement, anger) based on emotional analysis.	GPT-4 for categorization, manual review for accuracy.
Headline Category	The subject matter (e.g., politics, sports, entertainment, health), identified using topic modelling (LDA) for topical classification.	Topic modelling using Latent Dirichlet Allocation (LDA) [24]

The inclusion of emotion, arousal, and headline categories allows for deeper insights into how sensationalism operates within different contexts and across various emotional triggers.

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

Dataset Characteristics

The SENS-HEAD dataset offers a balanced mix of sensational and neutral headlines, crucial for training models capable of unbiased predictions. The distribution ensures that the models can effectively handle both extremes of sensationalism. Headlines span a wide array of topics, including politics, health, sports, and entertainment. This diversity ensures that the dataset is representative of real-world media, making the findings applicable across multiple domains. By incorporating sources from tabloid-style outlets and traditional journalism, the dataset captures a variety of writing styles and sensationalism tactics, providing a rich training ground for machine learning models.

Steps were taken to reduce potential biases in source selection and annotation. For instance, clickbait and non-clickbait sources were carefully curated to represent different journalistic practices, ensuring that the dataset does not overly favour one style. The current dataset focuses on English-language headlines. Future iterations could include multilingual data to account for cultural differences in sensationalism, as well as multimedia elements (e.g., images, videos) to capture non-textual cues that amplify sensationalism.

Necessity of the SENS-HEAD Dataset

Existing datasets, such as the News Clickbait Dataset [11] and similar resources, provide valuable insights into clickbait detection but fail to address sensationalism as a distinct phenomenon. Clickbait primarily focuses on misleading cues designed to drive clicks, often overlooking the emotional arousal and linguistic exaggeration central to sensationalism. Sensationalism goes beyond click-maximization strategies to evoke deeper psychological responses, shaping public perception and trust in media. Most available datasets have a limited scope, emphasizing clickbait without explicitly labelling or analyzing sensationalism. Additionally, they lack detailed annotations for emotional intensity and arousal categories, which are crucial for a nuanced understanding of sensationalism. Few datasets provide features suitable for explainable AI approaches, such as detailed reasoning for label assignments or interpretable linguistic features.

The SENS-HEAD dataset was specifically created to address these gaps and provide a comprehensive resource for sensationalism detection. Unlike existing datasets, SENS-HEAD focuses on emotional and linguistic patterns by including fields such as Emotion, Arousal, and Sensation Reason, enabling a deeper exploration of the psychological triggers used in sensational headlines. Furthermore, the dataset supports explainable AI approaches with features like Sensation Reason and detailed annotations, enhancing the interpretability and trustworthiness of models built using this data. Its broad coverage across diverse domains, including politics, health, entertainment, and sports, makes it more representative of real-world media compared to other domain-specific datasets. Additionally, SENS-HEAD ensures a balanced representation of sensational and neutral headlines, promoting unbiased training and robust model performance across different classes. By addressing these limitations, the SENS-HEAD dataset offers a unique and essential resource for advancing the study of sensationalism and bridging the gap between clickbait detection and broader emotional and linguistic analysis in news headlines.

Feature Extraction

Sensational headlines are often crafted to evoke strong emotional responses or curiosity, using exaggerated language, informal tones, and certain types of punctuation. By focusing on these linguistic elements, we can develop models capable of identifying sensational content, helping to address issues related to misinformation, clickbait, and ethical news dissemination. In this study, feature engineering plays a critical role in detecting sensationalism in news headlines by extracting a variety of linguistic, syntactic, and semantic features as depicted in figure 2.

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

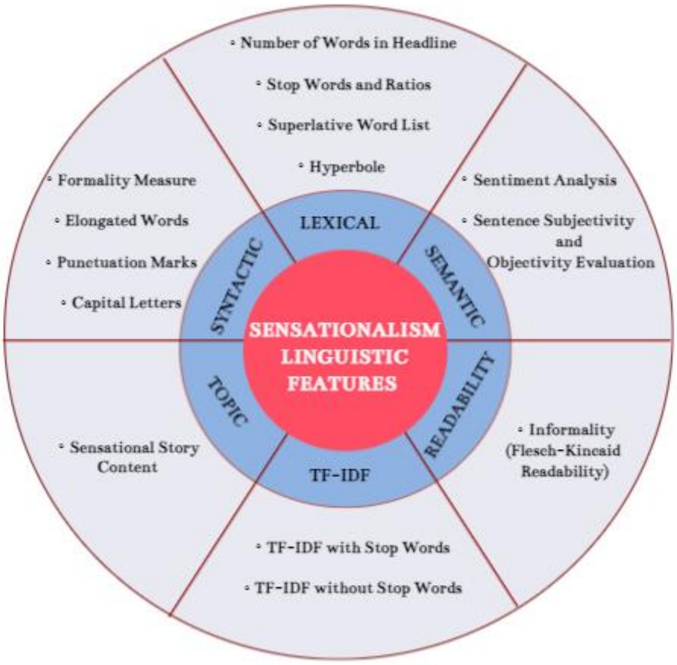

Fig.2. Sensationalism Feature Wheel

These features are designed to capture different aspects of the language used in headlines, ranging from simple lexical patterns to more complex syntactic structures and emotional tones. By analyzing these features, the model can differentiate between sensational headlines and more neutral ones. The following sub-section breaks down the different categories of features used in the study and explaining their significance in detecting sensationalism.

Lexical Features

Lexical features focus on individual words and phrases within the headlines, providing insight into the language's simplicity or complexity, as well as the presence of exaggerated or sensational terms.

- **Number of Words in Headline**: The total word count in each headline was extracted to determine the length. The total word count is important because shorter, punchier headlines tend to be more sensational. For example, *"Shocking! You Won't Believe What Happened!"* contains fewer words but is highly attention-grabbing compared to a more neutral, fact-based headline like *"Survey Results Released on Public Opinion."* The word count was calculated by tokenizing the headline and counting the tokens.

- **Stop Words and Ratios**: The ratio of stop words (e.g., "the," "is") to content words (e.g., nouns, verbs) helps identify whether the language is more informal or simplified. Sensational headlines often use simpler language to appeal to a wider audience. For example, a headline like *"This Is What You Need to Know"* uses more stop words than a more complex version, such as *"Key Findings in Latest Research Report."* A predefined list of stop words was used for this purpose.

- **Superlative Words**: Superlative adjectives and adverbs (e.g., "best," "most amazing") are frequently used in sensational headlines to exaggerate the importance or appeal of an event or product. For example, *"The Best Strategy You'll Ever Find!"* is clearly designed to entice clicks through exaggeration, compared to a more factual version, *"Effective Strategies for Time*

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

Management." The Superlative Lexicon consisting of 594 words, was created. These words were extracted from multiple sources, including internet lists and five key NLTK corpora: Brown, Gutenberg, Movie Reviews, Reuters, and Webtext.

- **Hyperbole (Extreme Case Formulation)**: Hyperbole is another common tactic used to sensationalize headlines, often through extreme language. For example, *"The Worst Day in History!"* uses hyperbole to amplify the emotional impact, compared to *"A Challenging Day for Many."* A custom lexicon of 160 hyperbolic terms was created to detect exaggerated or extreme language, which is commonly found in sensational headlines. These terms were extracted using regular expressions to identify exaggerated claims.

Syntactic Features

Syntactic features focus on the structure of sentences and the use of parts of speech, which can signal the formality or informality of the language. Sensational headlines often rely on informal, simple syntactic structures.

- **Formality Measure (F-measure)**: This feature helps distinguish between formal and informal language by analyzing the frequency of different parts of speech. For example, formal headlines like *"Government Enacts Comprehensive Legislation"* typically contain more nouns, adjectives, and prepositions. In contrast, informal, sensational headlines like *"You Won't Believe What They Did!"* rely more on pronouns and verbs, making the language more dynamic and casual. This feature measures the degree of formality in a headline by calculating the frequency of different parts of speech. Formal language typically uses more nouns, adjectives, prepositions, and articles, while informal language relies on pronouns, verbs, adverbs, and interjections. The F-measure is calculated using the Heylighen and Dewaele [25] formula which is given as follows:

$$Formality = \frac{(nounfreq + adjectivefreq + prepositionfreq + articlefreq) - (pronounfreq + verbfreq + interjectionfreq) + 100}{2}$$

where,

- *Noun Frequency (nounfreq):* Higher noun frequencies are associated with more formal language as nouns provide specific information and descriptions.
- *Adjective Frequency (adjectivefreq):* Higher adjective frequencies add more detail and specificity, which is also characteristic of formal language.
- *Preposition Frequency (prepositionfreq):* Prepositions contribute to sentence structure and complexity, making the language more formal.
- *Article Frequency (articlefreq):* Articles (e.g., "the," "a," "an") mark the definiteness of noun phrases and are common in formal writing.

The negative components lower the formality score:

- *Pronoun Frequency (pronounfreq):* Pronouns are more common in informal language as they replace specific nouns, making sentences less detailed.
- *Verb Frequency (verbfreq):* Simple verbs are typically used in informal language as they make sentences more dynamic and concise.
- *Adverb Frequency (adverbfreq):* Adverbs often modify verbs and express emotions, contributing to a conversational tone.
- *Interjection Frequency (interjectionfreq):* Interjections (e.g., "Wow!", "Oh no!") are common in spoken, informal language.

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

The inclusion of the constant 100 and division by 2 ensures that the score is within a reasonable range, typically between 0 and 100, making it easier to interpret and compare across texts.

- **Elongated Words (EW)**: Elongated words, such as *"soooo amazing,"* are often used to convey emphasis in sensational headlines. They indicate informal language designed to amplify emotions. For instance, *"This Is Sooo Cool!"* is more sensational than *"This Is Cool."*

- **Punctuation Marks (PUNC)**: Sensational headlines frequently employ punctuation to dramatize their message. For example, exclamation marks or ellipses are often used to create suspense or urgency, as in *"Breaking News: Market Crash Looms!"* or *"You Won't Believe What Happened Next…"*

- **Capital Letters (CL)**: Headlines that use capital letters to emphasize certain words often appear more sensational. For instance, *"INCREDIBLE OFFER!"* grabs attention far more aggressively than *"Incredible Offer."*

Semantic Features

Semantic features focus on the meaning and emotional tone of the headlines. Sensational headlines often convey strong emotions or subjective opinions to provoke reactions from readers.

- **Sentiment Analysis**: Sentiment scores (positive, negative, neutral, and compound) were calculated to assess the emotional tone of each headline. Headlines that evoke strong positive or negative emotions are more likely to be sensational. For example, a headline like *"Horrific Tragedy Strikes!"* contains highly negative sentiment, compared to a neutral, fact-based headline like *"Accident Reported on Highway."* Sentiment analysis tools were used to extract this information.

- **Sentence Subjectivity and Objectivity Evaluation**: This feature evaluates whether a headline is subjective (opinion-based) or objective (fact-based). Sensational headlines are typically more subjective, aiming to provoke emotional responses. For example, *"The Best Product You Must Try!"* is subjective, while *"New Product Released in Stores"* is objective.

Readability Features

Readability features measure the linguistic complexity of headlines. Sensational headlines are often simpler to make them more accessible to a broader audience.

- **Informality (Flesch-Kincaid Readability) [26]**: This feature measures the degree of informality in the headline's language. Sensational headlines typically aim for a conversational tone, such as "*Find Out How This Hack Can Save You Time!"* compared to *"Time-Saving Techniques Explained."*

TF-IDF Features

TF-IDF features [27] measure the importance of individual terms in the context of the entire corpus of headlines, helping to identify sensational terms that appear frequently in sensational headlines but rarely in neutral ones.

- **TF-IDF with Stop Words**: By including stop words, this feature helps capture frequently used phrases that contribute to sensationalism, such as *"the best"* or *"you won't believe."*

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

- **TF-IDF without Stop Words**: Excluding stop words refines the analysis by focusing on content words that are more meaningful and indicative of sensationalism.

Topic and Content-Based Features
Topic-based features help identify whether the content of the headline relates to sensational topics like crime, disasters, or scandals.

- **Sensational Story Content**: Headlines related to crime, disasters, and scandals are often sensational by nature. For example, *"Massive Earthquake Destroys City!"* is more likely to be sensational compared to *"City Experiences Minor Earthquake."*

This comprehensive set of features was extracted and processed using various NLP techniques and tools, which are outlined in the following table along with the specific tools used for each feature:

Table 2. Summary of Feature Engineering

Feature Category	**Linguistic Feature**	**Description**	**Tool Used**
Lexical Features	Number of Words in Headline	Total word count in each headline. Indicates headline length and its impact on sensationalism.	Python string operations, spaCy
	Stop Words and Ratios	Measures the number and ratio of stop words to content words.	spaCy built-in stop words list
	Superlative Word List	Tracks superlative adjectives/adverbs to detect exaggeration. 594 words from NLTK corpora and internet sources.	Manually crafted lexicon
	Hyperbole (Extreme Case Formulation)	Detects exaggerated claims using a manually crafted lexicon of 160 hyperbolic terms.	Manually crafted lexicon
Syntactic Features	Formality Measure (F-measure)	Measures formality using noun/adjective/preposition/article frequency versus pronoun/verb/adverb/interjection frequency.	Custom formula using spaCy POS tags
	Elongated Words (EW)	Tracks elongated words (e.g., "soooo amazing"), often seen in sensational language.	Regular expressions
	Punctuation Marks (PUNC)	Analyzes specific punctuation (e.g., exclamation marks, ellipses) and informal marks.	Regular expressions, Python string operations
	Capital Letters (CL)	Detects capital letters used for emphasis in headlines.	Regular expressions
Semantic Features	Sentiment Analysis	Calculates positive, negative, neutral, and compound sentiment to capture emotional tone.	VADER
	Sentence Subjectivity and Objectivity Evaluation	Determines whether the headline is subjective (opinionated) or objective (factual).	TextBlob

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

Readability Feature	Informality (Flesch-Kincaid Readability)	Measures the informality of headlines using readability metrics.	textstat library
TF-IDF Features	TF-IDF with Stop Words	Measures word importance with stop words included.	scikit-learn TfidfVectorizer
	TF-IDF without Stop Words	Measures word importance with stop words excluded.	scikit-learn TfidfVectorizer
Topic and Content-Based Feature	Sensational Story Content	Identifies sensational topics using keyword matching and WordNet.	WordNet, Regular expressions

Model Training

The model training process was designed to be an incremental, multi-stage approach to ensure that the most significant features were identified and used effectively. Each stage aimed to improve the model's ability to detect sensationalism in headlines by refining the feature set and employing ensemble machine learning models to enhance classification accuracy. The process involved three key stages: assessing the contribution of individual features, determining an optimal threshold using Principal Component Analysis (PCA) [28], and training various ensemble models on all features to identify the best-performing algorithm.

Stage 1: Train with Individual Features

The first stage focused on training the model using each feature individually. By evaluating each feature in isolation, we were able to assess its individual contribution to detecting sensationalism. This step helps to identify the most significant features that impact model performance before moving to a more complex, multi-feature approach. Training with individual features allows us to gain insights into which features are most predictive of sensationalism. Some features, such as sentiment analysis and punctuation usage, may have a more direct impact on the classification of sensational headlines than others like the number of words or readability scores. Evaluating features in isolation ensures that we can identify and prioritize features that provide the most significant improvement in model performance.

Each of the following features was trained individually to assess its impact on the overall model's ability to detect sensationalism. The Number of Words feature was used to evaluate the effect of headline length, while the Number of Stop Words captured the presence of common words that typically contribute to simpler, more accessible language. The Ratio of Stop Words to Content Words was calculated to measure the balance between stop words and meaningful content words, providing insight into the complexity of the headline's language structure. The Flesch-Kincaid Readability score was used to assess the linguistic complexity of the headlines, and the Subjectivity and Objectivity feature evaluated whether the headline appealed to emotions (subjective) or reported factual information (objective). To further capture emotional tone, Sentiment Analysis was performed, generating scores for negative, neutral, positive, and compound sentiments. Additionally, the Elongated Words feature tracked exaggerated word forms (such as "sooo") commonly used in sensational content. The Punctuation feature analyzed the presence of dramatic punctuation marks, including exclamation marks, question marks, ellipses, and currency symbols, which are often employed to amplify the headline's impact. Lastly, TF-IDF with Stop Words measured term importance by including stop words, while TF-IDF without Stop Words refined this measure by excluding them, helping to highlight the most meaningful words related to sensationalism.

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

This phase was crucial in isolating which features had the greatest impact on classification accuracy. By training with these individual features, we developed a clear understanding of which attributes of headlines were most likely to indicate sensationalism.

Stage 2: Calculate Optimal Threshold for XGBoost

Once the most impactful individual features were identified, the next step was to enhance the model by using Principal Component Analysis (PCA). PCA was applied specifically to the emotion data to reduce the dimensionality and focus on the most significant components. PCA helps in simplifying the dataset by transforming the emotion features into principal components. This transformation retains the most critical information while reducing complexity, making it easier for the model to interpret and classify the data. PCA is especially useful when dealing with high-dimensional data like emotion scores, where there might be correlations between variables. The steps in Stage 2 included:

- **Principal Component Analysis (PCA)**: PCA was applied to the emotion column to identify the most significant components that contributed to classifying headlines as sensational or non-sensational. This step helped reduce noise in the data by focusing on the key features that had the most impact on the classification.
- **Combine Sensation and Arousal Scores**: After applying PCA, the sensation score and arousal score were combined with the principal components. These scores represent the intensity of sensationalism, and the emotional arousal evoked by the headline, respectively. By combining these scores, we created a more comprehensive feature set that captured both the sensational and emotional aspects of headlines.
- **XGBoost Baseline Model**: The combined features (sensation, arousal, and PCA components) were then used to train an initial XGBoost classifier [29]. XGBoost is a gradient boosting algorithm known for its robustness and efficiency in classification tasks. The optimal threshold for classification was determined based on this model's performance, allowing the model to make more accurate predictions when classifying headlines as sensational or non-sensational.

Stage 3: Train Final Models

In the final stage of the training process, the focus was on combining the most relevant features into a comprehensive feature set to maximize model performance while reducing computational complexity. Before concatenating all features for model training, we optimized feature selection by applying a variance threshold of 0.001 to remove constant or near-constant features. This step effectively reduced the feature set from 52,190 to 1,236 features, eliminating those that contributed little to the model's predictive power. This process not only minimized the risks of overfitting but also conserved computational resources and improved model interpretability by narrowing the focus to more meaningful features. By concatenating all selected features (except for TF-IDF without stop words), we created a complete and representative dataset for training. The objective was to leverage this enhanced feature set to explore and compare the performance of multiple ensemble machine learning models, each offering unique strengths in handling complex data structures.

Ensemble learning, a powerful technique that combines predictions from multiple models, was central to this stage. By incorporating models such as XGBoost, CATBoost, AdaBoost, and Random Forest, we aimed to boost predictive accuracy through the complementary nature of these algorithms. Ensemble methods are known for their ability to reduce overfitting and enhance generalization by combining the strengths of individual models into a more robust final prediction. This diverse set of models allowed us to exploit various aspects of the data, including linguistic, syntactic, and semantic features, ultimately producing more reliable results in detecting sensationalism in news headlines. Table 3 gives the summary of the models trained and their specific configurations:

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

Table 3: Summary of ensemble learning models trained with specific features and configurations for sensationalism detection.

Model	Description
XGBoost	Trained using all features, XGBoost is renowned for its ability to handle imbalanced data while offering fast and accurate classification results.
XGBoost with Superlative Adjective Words List	This variant included the superlative adjectives feature to evaluate the impact of exaggerated adjectives on detecting sensationalism in headlines.
XGBoost with Threshold	The optimal threshold, determined in Stage 2, was applied to XGBoost, improving its ability to classify borderline cases of sensationalism more effectively.
XGBoost with Superlative Adjective Words List and Threshold	A combination of the superlative adjectives feature and the optimal threshold, this variant offers a more fine-tuned approach to headline classification.
AdaBoost	An ensemble model trained on all features, AdaBoost combines weak classifiers to build a strong classifier. It focuses on correcting the mistakes of previous models, improving overall performance in detecting sensational headlines.
CATBoost	Another gradient boosting algorithm, CATBoost is particularly efficient for handling categorical data and is a faster alternative to XGBoost without sacrificing performance. Trained on all features.
Random Forest	A traditional ensemble method, Random Forest builds multiple decision trees and averages their predictions to improve classification accuracy while reducing the risk of overfitting. It was trained using all available features.
CATBoost on Test Set	In the final evaluation step, CATBoost was tested on the test set to measure its performance in detecting sensationalism in unseen data, providing insights into the model's generalization capabilities.

By following this methodical approach, we ensured that the model was not only capable of accurately detecting sensationalism but also robust against noise and less impactful features. Each model was evaluated for its ability to classify sensational headlines accurately, and the most effective model was selected for further analysis.

Model Implementation and Evaluation

In this section, we detail the steps taken for model implementation and evaluation, aiming to detect sensationalism in news headlines. We selected XGBoost as the baseline model due to its robustness and efficiency in handling structured data, making it ideal for our dataset. To compare performance, we also employed other ensemble models such as AdaBoost, CATBoost, and Random Forest. These models offer distinct advantages in handling complex, high-dimensional datasets, making them valuable for benchmarking. We chose not to experiment with contextual embeddings and stacked models in this study to prioritize explainability, computational feasibility, and alignment with our research objectives. Advanced techniques like BERT and stacked models, while powerful, often reduce interpretability, which is critical for practical applications in media ethics and policymaking.

Given the computational constraints and the high dimensionality of the feature space, we opted for Random Search instead of Grid Search for hyperparameter tuning. To ensure consistent performance across different subsets of the data, we used 5-fold cross-validation (n_splits = 5). This approach divides

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

the data into five folds and trains the model on different combinations of these folds, enhancing the reliability of the model's performance.

The F1-score was chosen as the primary scoring metric because it balances precision and recall, making it particularly useful for datasets with class imbalances. This metric provides a comprehensive evaluation of the model's ability to detect both sensational and non-sensational headlines accurately. The model implementation followed a structured approach. First, we split the dataset into training and validation sets to assess the models' performance on unseen data. Each model was then trained using the optimized feature set and hyperparameters, with evaluation metrics such as precision, recall, and F1-score calculated for both the sensational and non-sensational classes. Finally, the best-performing model was further tested on a separate test set to assess its generalization capabilities, ensuring it performed well beyond the validation data.

Table 4 below summarizes the key hyperparameters used across the four ensemble machine learning models. While these models share common parameter types such as the number of estimators, learning rate, tree depth, and feature subsampling, each model's implementation varies slightly. For example, while XGBoost and CATBoost use regularization and subsampling techniques, Random Forest relies on bootstrap sampling and does not use a learning rate. AdaBoost, on the other hand, uses different algorithms (SAMME, SAMME.R) to adjust its boosting approach. This table highlights the similarities and differences in parameter configurations across the models, illustrating how each model was fine-tuned to address performance and computational efficiency in our experiments.

Table 4: Hyperparameter tuning parameters for ensemble models.

Parameter Type	**XGBoost**	**AdaBoost**	**CATBoost**	**Random Forest**
Number of Estimators	'n_estimators': [100, 300, 500]	'n_estimators': [100, 300, 500]	'iterations': [100, 300, 500]	'n_estimators': [100, 300, 500]
Learning Rate	'learning_rate': [0.01, 0.1, 0.3]	'learning_rate': [0.01, 0.1, 0.3]	'learning_rate': [0.01, 0.1, 0.3]	N/A
Tree Depth	'max_depth': [3, 6, 9]	'base_estimator__max_depth': [3, 6, 9]	'depth': [3, 6, 9]	'max_depth': [3, 6, 9]
Min Samples per Split	'min_child_weight': [1, 3]	'base_estimator__min_weight_fraction_leaf': [0, 0.1]	'l2_leaf_reg': [1, 3]	'min_samples_leaf': [1, 3]
Feature Subsampling	'colsample_bytree': [0.8, 1.0]	'base_estimator__max_features': [0.8, 1.0]	'colsample_bylevel': [0.8, 1.0]	'max_features': [0.8, 1.0]
Row Subsampling	'subsample': [0.8, 1.0]	N/A	'subsample': [0.8, 1.0]	'bootstrap': [True, False]
Regularization	'gamma': [0, 0.1]	N/A	'random_strength': [0, 0.1]	'min_impurity_decrease': [0, 0.1]
Model-Specific Parameters	N/A	'algorithm': ['SAMME', 'SAMME.R']	N/A	N/A

Model Performance and Evaluation

To evaluate the effectiveness of each model, we analyzed the performance metrics using the validation set. The metrics included precision, recall, F1-score, and support for both sensational and non-sensational headline classes. These metrics offer a comprehensive assessment of how well the models performed across different classification tasks, balancing the need for precision and recall in a dataset

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

characterized by class imbalances [30]. The performance results for each model varied slightly, depending on the feature set and configurations:

- XGBoost with Superlative Adjective Words List demonstrated balanced performance, with similar F1-scores for both sensational and non-sensational classes, showing it captured the nuances of exaggerated language effectively.
- XGBoost with Threshold showed a slightly stronger bias toward detecting sensationalism, reflecting a more cautious approach to classifying sensational content. This suggests it was more effective in identifying sensational headlines but with a trade-off in detecting non-sensational ones.
- AdaBoost performed well, excelling at correcting misclassified instances, particularly in the sensational class, making it effective in handling difficult-to-predict headlines.
- CATBoost maintained consistent F1-scores across both classes, indicating that it generalized well across the dataset and effectively handled the complexities of sensationalism detection.
- Random Forest showed slightly lower performance, suggesting potential issues with either overfitting or underfitting, indicating that it may not have been able to capture the complexity of the dataset as effectively as the other models.

Table 5 shows the performance metrics for XGBoost with Superlative Adjective Words List. The XGBoost model with the Superlative Adjective Words List performed consistently across both classes, demonstrating a slight advantage in detecting non-sensational headlines. The addition of the Superlative Adjective Words List provided the model with improved sensitivity to exaggerated language, which is often indicative of sensationalism.

Table 5: XGBoost performance with Superlative Adjective Words List

Metric	Non-Sensation	Sensation
Precision	0.68	0.64
Recall	0.67	0.65
F1-Score	0.68	0.64
Support	3,223	2,862

The ROC-AUC curve (Figure 3) demonstrates the XGBoost model's ability to distinguish between sensational and non-sensational headlines, with an AUC of 0.72 indicating moderate performance. The model strikes a balance between true positive rate and false positive rate but has room for improvement in specificity and sensitivity.

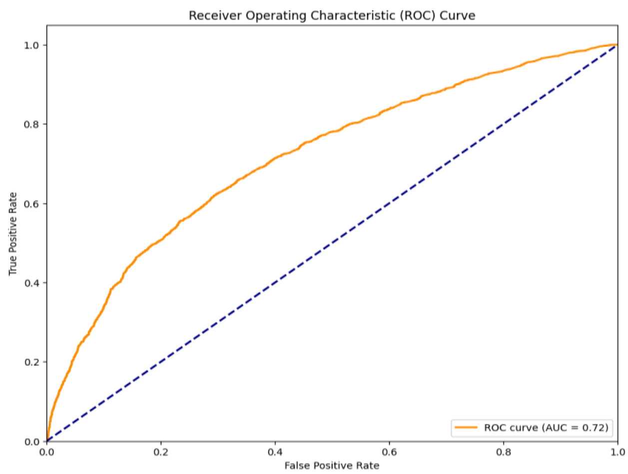

Fig.3. ROC-AUC curve for XGBoost with Superlative Adjective Words List

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

The Confusion Matrix (Figure 4) provides a breakdown of predictions. The model correctly classified 2,162 non-sensational headlines and 1,859 sensational headlines. However, 1,061 non-sensational headlines were incorrectly classified as sensational (false positives), while 1,003 sensational headlines were missed and classified as non-sensational (false negatives). This balance suggests the model performs reasonably well but could be improved in terms of both precision and recall.

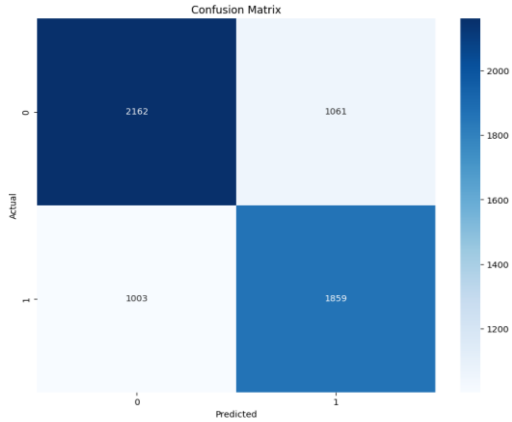

Fig.4. Confusion Matrix for XGBoost with Superlative Adjective Words List

Applying an optimal threshold significantly improved the model's ability to detect sensational headlines, reflected in the higher recall for the sensation class. However, there was a trade-off with precision for the non-sensation class, as the model became more cautious, favouring sensation detection over non-sensation classification. This resulted in a balanced performance for both classes, but with a slight edge toward detecting sensational content as shown in table 6.

Table 6: XGBoost performance with Threshold

Metric	**Non-Sensation**	**Sensation**
Precision	0.70	0.61
Recall	0.59	0.72
F1-Score	0.64	0.66
Support	3,223	2,862

The Confusion Matrix (Figure 5) shows the performance of the XGBoost model with Threshold. The model correctly classified 1,897 non-sensational headlines and 2,056 sensational headlines. However, it misclassified 1,326 non-sensational headlines as sensational (false positives) and 806 sensational headlines as non-sensational (false negatives). This indicates that while the model is relatively effective at detecting sensational content, it still misclassifies a significant portion of non-sensational headlines, suggesting room for improvement in balancing precision and recall.

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

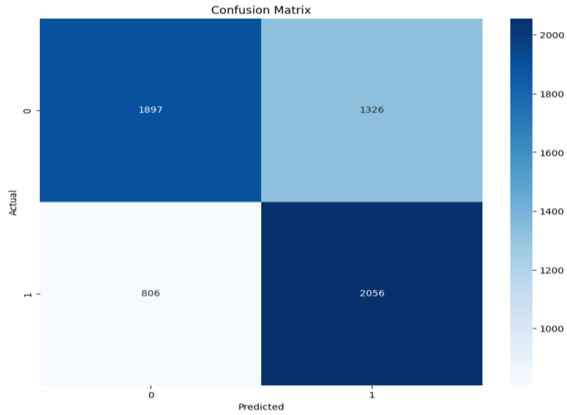

Fig.5. Confusion Matrix for XGBoost with Threshold

The ROC-AUC curve for XGBoost with Threshold in figure 6, shows the model's ability to distinguish between sensational and non-sensational headlines, with an AUC of 0.71.

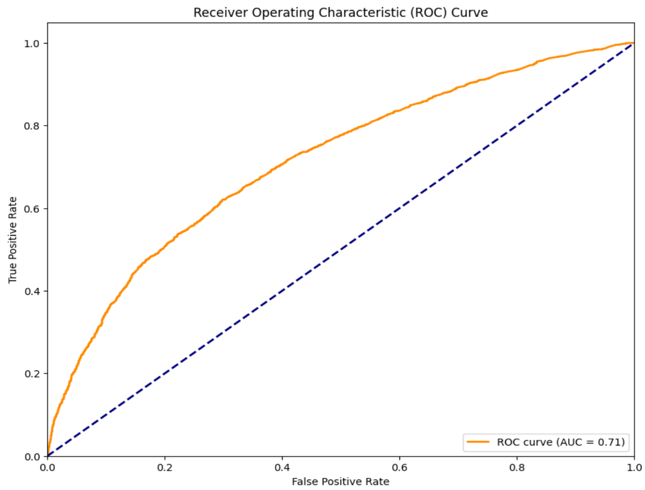

Fig.6. ROC-AUC curve for XGBoost with Threshold

Combining the Superlative Adjective Words List with the optimal threshold created a fine-tuned model that performed similarly to the threshold-only model as shown in table 7. The addition of superlative adjectives slightly improved the model's ability to detect exaggerated language, while maintaining a balance between sensational and non-sensational headline detection.

Table 7: XGBoost performance with Superlative Adjective Words List and Threshold

Metric	Non-Sensation	Sensation
Precision	0.70	0.61
Recall	0.59	0.72
F1-Score	0.64	0.66
Support	3,223	2,862

The Confusion Matrix (Figure 7) for the XGBoost model with Superlative Adjective Words List and Threshold shows that the model correctly classified 1,888 non-sensational headlines and 2,071 sensational headlines. However, it misclassified 1,335 non-sensational headlines as sensational (false positives) and 791 sensational headlines as non-sensational (false negatives). This matrix suggests that

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

while the model is relatively good at detecting sensational content, the inclusion of the superlative adjective list improves its performance slightly, though further refinements are needed to reduce misclassifications, particularly false positives.

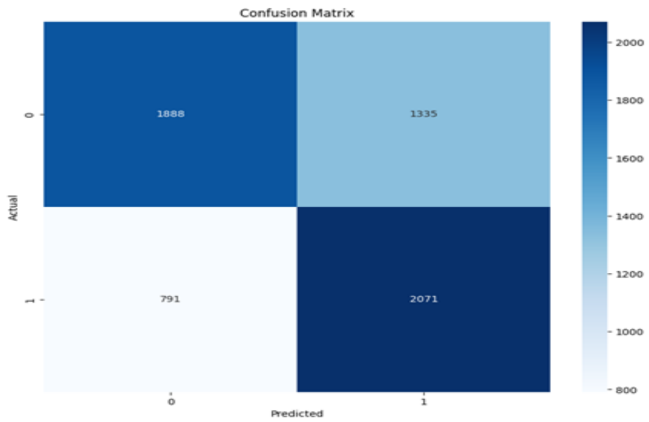

Fig.7. Confusion Matrix for XGBoost with Superlative Adjective Words List and Threshold

The ROC-AUC Curve (Figure 8) for the XGBoost model with Superlative Adjective Words List and Threshold shows an AUC of 0.72, which is the same as the model using the Superlative Adjective Words List alone. While this does not indicate any significant improvement in performance when combining the two techniques, it still reflects moderate effectiveness in distinguishing between sensational and non-sensational headlines. The model's sensitivity to exaggerated language is maintained, but the combination of the two features does not provide additional benefit over using the superlative adjective list independently. This suggests that further adjustments might be necessary to fully leverage both features for better classification.

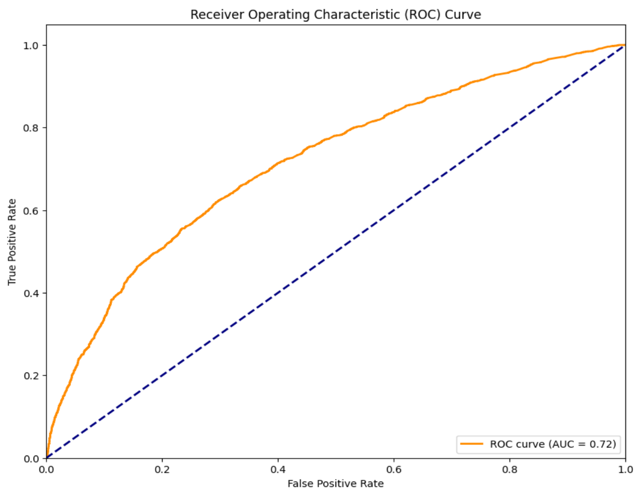

Fig.8. ROC-AUC curve for XGBoost with Superlative Adjective Words List and Threshold

The table 8 compares the performance of AdaBoost, CATBoost, and Random Forest in terms of Precision, Recall, and F1-Score for both non-sensational and sensational classes. The results in the table reveal distinct performance patterns across the models. AdaBoost demonstrates a balanced trade-off between precision and recall for both non-sensational (Precision: 0.68, Recall: 0.65) and sensational headlines (Precision: 0.62, Recall: 0.66), suggesting it handles both classes fairly evenly. CATBoost achieves slightly higher overall consistency, maintaining equal precision and recall for both non-

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index
sensational and sensational headlines (0.68 and 0.64 respectively), making it the most balanced model. In contrast, Random Forest shows a notable disparity: higher precision for non-sensational headlines (0.69) but much lower precision for sensational headlines (0.59), though it compensates with higher recall for sensational content (0.70). This suggests Random Forest may favour detecting sensationalism at the cost of misclassifying non-sensational content.

Table 8: Performance Metrics for AdaBoost, CATBoost, and Random Forest

Model	Metric	Non-Sensation	Sensation
AdaBoost	**Precision**	0.68	0.62
	Recall	0.65	0.66
	F1-Score	0.67	0.64
	Support	3,223	2,862
CATBoost	**Precision**	0.68	0.64
	Recall	0.68	0.64
	F1-Score	0.68	0.64
	Support	3,223	2,862
Random Forest	**Precision**	0.69	0.59
	Recall	0.57	0.70
	F1-Score	0.62	0.64
	Support	3,223	2,862

Figure 9 depicts the boxplot comparing Precision, Recall, and F1-Score for the three models (AdaBoost, CATBoost, and Random Forest).

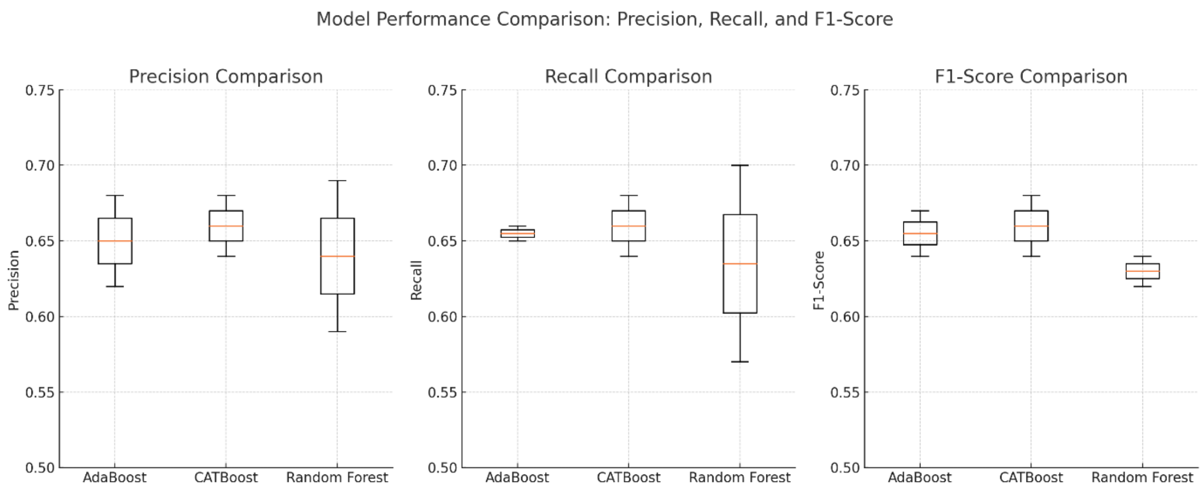

Fig.9. Box plot comparison of Precision, Recall, and F1-Score across AdaBoost, CATBoost, and Random Forest

The box plots provide a clear visual comparison of the performance of the three models — AdaBoost, CATBoost, and Random Forest — in terms of precision, recall, and F1-score. Each box plot displays the interquartile range (IQR), with the median represented by the line inside each box, and the whiskers extending to the range of the data, excluding outliers.

For precision, both AdaBoost and CATBoost show relatively narrow and consistent distributions, with similar medians, indicating stable precision across both models. Random Forest, on the other hand,

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

shows more variability in precision, as its wider IQR suggests fluctuations in how accurately it detects the correct class (sensational vs non-sensational). In the case of recall, Random Forest displays significantly more variability, with a wider range in performance. This implies that the model is less consistent in correctly identifying true positives, especially for one class. In contrast, CATBoost and AdaBoost exhibit tighter and more stable recall distributions, indicating that they are more reliable in identifying true positive instances in both classes.

The F1-score plot similarly shows that Random Forest experiences more variability compared to AdaBoost and CATBoost. While AdaBoost and CATBoost demonstrate consistent performance with relatively high medians, Random Forest's wider IQR suggests that it may struggle with balancing precision and recall effectively. This inconsistency might be due to overfitting or underfitting when dealing with certain feature combinations. Overall, the box plots highlight CATBoost and AdaBoost as the more stable and balanced models in terms of precision, recall, and F1-score, while Random Forest demonstrates wider variability, which may indicate challenges in consistently handling the data or target classes.

Comparative Analysis with Prior Approaches

No prior work has directly addressed sensationalism detection in news headlines using the specific combination of linguistic features and explainable ensemble models as proposed in this study. However, the nearest works in the field provide valuable comparisons for contextualizing our contributions.

Zhang and Kejriwal [22] relied on linguistic features such as part-of-speech tags, punctuation frequency, sentence length, and polarity/subjectivity scores, utilizing a Support Vector Machine (SVM) classifier for sensationalism detection. Their study focused on U.S. political news headlines from 2017 and 2019, sourced from MediaBiasFactCheck.com, with an emphasis on concept drift. While their models achieved a precision of 0.91 and recall of 0.93 on 2017 data, their performance declined significantly on 2019 data, highlighting challenges in generalizability due to evolving linguistic trends.

In contrast, our study introduces a novel approach, leveraging advanced ensemble models like XGBoost, CATBoost, Random Forest, and AdaBoost, enriched with domain-specific features such as Sensation Score, Emotion, and Arousal. Using the SENS-HEAD dataset, which spans diverse domains beyond politics, including health, entertainment, and sports, our models demonstrated higher stability and broader applicability. CATBoost emerged as the most consistent model, achieving an F1-score of 0.68 for both sensational and non-sensational classes, while XGBoost with optimized thresholds reached an F1-score of 0.66 for the sensational class, indicating robust performance across classes.

Similarly, González Esparza et al. [31] focused on Spanish health-related headlines using deep learning models like Multilingual BERT, BETO, and XLM-RoBERTa. Their models achieved impressive results, with F1-scores reaching 0.94, and XLM-RoBERTa showing strong recall (96%). However, their approach was limited to a single domain and language, and it lacked the interpretability mechanisms central to our SHAP-based analysis. Our framework balances performance and explainability, making it more adaptable for real-world applications across multiple domains.

The integration of SHAP analysis in our study provides critical transparency, enabling the identification of influential features such as headline length, stop word frequency, and sentiment polarity. This approach not only enhances model interpretability but also ensures practical applicability for content moderation and ethical journalism. This comparative analysis underscores that while prior works have made strides in related areas, our study uniquely combines linguistic features, explainable models, and broad domain coverage to advance sensationalism detection. Future work will focus on incorporating

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

contextual embeddings and dynamic retraining to address evolving language trends and concept drift more effectively.

Explainability Techniques and SHAP Analysis

In machine learning, especially with complex models like ensemble methods, explainability is critical for understanding how models make predictions [32]. One of the most effective explainability techniques is the SHapley Additive exPlanations (SHAP) [6] framework, which provides global and local explanations for individual predictions. SHAP assigns an importance value to each feature in a way that the sum of the feature contributions equals the difference between the actual and the baseline prediction. SHAP is grounded in cooperative game theory and uses Shapley values to explain the contribution of each feature. This is important in applications such as sensationalism detection, where understanding why the model identifies a headline as sensational can inform future improvements in media ethics and content regulation. In this context, SHAP values provide transparency, helping to pinpoint which linguistic features most heavily influence the classification of headlines as sensational or non-sensational.

The SHAP analysis was conducted to interpret the machine learning model's predictions and assess the significance of different features in detecting sensationalism in news headlines. SHAP is particularly suited for explaining the contribution of various lexical, syntactic, and semantic features used in this study. The SHAP Feature Importance Summary in Figure 10 provides a clear visualization of the most influential features driving the model's predictions in detecting sensationalism.

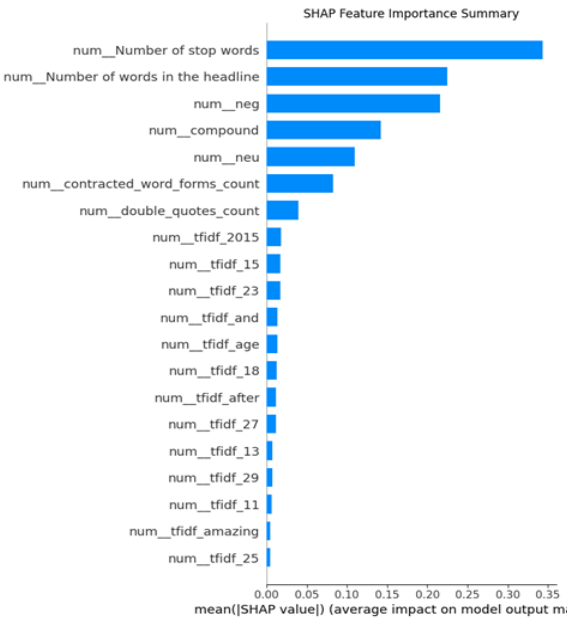

Fig. 10. SHAP Feature Importance Summary

The SHAP feature importance summary shows the average impact of each feature on the model's output (sensation detection). Here, the number of stop words and number of words in the headline emerged as the most influential features, suggesting that shorter, simpler headlines with fewer content words tend to be more sensational. Negative sentiment and compound sentiment scores also contributed significantly to the model's predictions, indicating that headlines with strong negative emotions or a mix of emotions were more likely to be classified as sensational. From this summary, we selected the top seven features—number of stop words, number of words in the headline, negative sentiment

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

(num_neg), compound sentiment (num_compound), neutral sentiment (num_neu), contracted word forms, and double quotes count—to create detailed SHAP dependence plots.

SHAP dependence plots help visualize the relationship between individual features and the model's output, allowing us to see how the feature's value affects the prediction. Importantly, these plots also reveal interactions between features that may not be obvious from feature importance alone. The SHAP Dependence Plot in Figure 11 highlights the relationship between the number of stop words and their impact on model predictions for sensationalism.

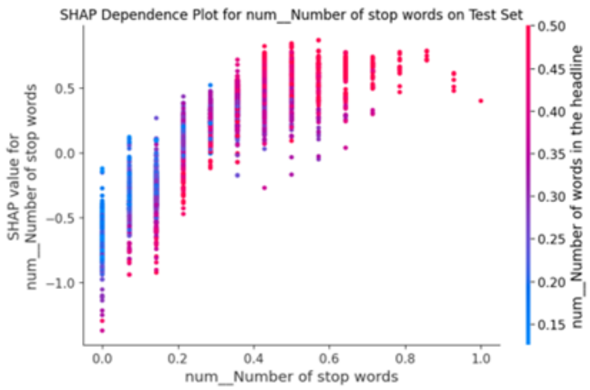

Fig. 11. SHAP Dependence Plot for Number of Stop Words on Test Set

This dependence plot highlights the relationship between the number of stop words in a headline and the SHAP value. The plot shows a positive correlation, where headlines with more stop words (above 0.2 on the x-axis) are more likely to be predicted as sensational. The breakthrough point occurs around 0.4, where the model becomes increasingly confident in predicting a headline as sensational. The colour gradient also shows how the number of words in the headline interacts with this feature.
The SHAP Dependence Plot in Figure 12 illustrates how the number of words in the headline influences the model's predictions for sensationalism.

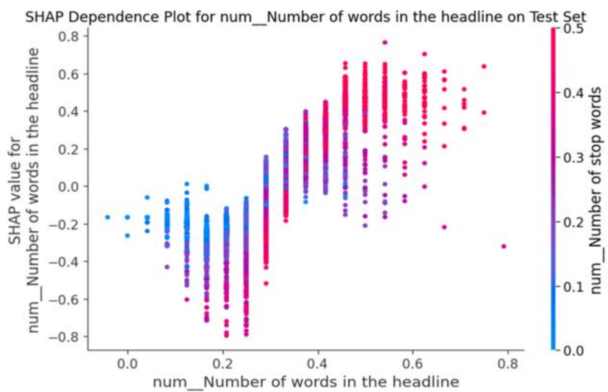

Fig.12. SHAP Dependence Plot for Number of Words in the Headline on Test Set

This plot demonstrates the influence of headline length (in terms of the number of words) on prediction outcomes. As the number of words increases, the model starts assigning higher SHAP values, especially in the range of 0.3-0.4, indicating a higher likelihood of the headline being classified as sensational. The interaction with stop words further emphasizes how shorter, pithy headlines with fewer content words tend to be more sensational.

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

The SHAP dependence plot in figure 13 highlights the relationship between compound sentiment scores and the model's predictions. Compound sentiment captures the overall emotional tone of the headline, ranging from extremely negative to extremely positive. The U-shaped curve indicates that both highly positive and highly negative headlines tend to be classified as sensational, with neutral headlines showing a lower likelihood of being classified this way. The interaction with the neutral sentiment scores (num_neu) further emphasizes that headlines with stronger emotional tones, whether positive or negative, are more likely to be considered sensational.

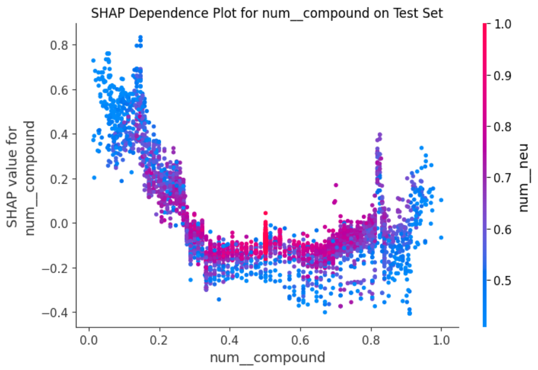

Fig. 13. SHAP Dependence Plot for Compound Sentiment on Test Set

The SHAP dependence plot for negative sentiment (num_neg) in figure 14 illustrates that as negative sentiment scores increase, the SHAP value also rises, indicating that more negative headlines are more likely to be classified as sensational. The color gradient representing neutral sentiment (num_neu) reinforces this, showing that highly negative headlines are less neutral and, thus, more likely to evoke stronger emotional reactions, contributing to their sensational classification.

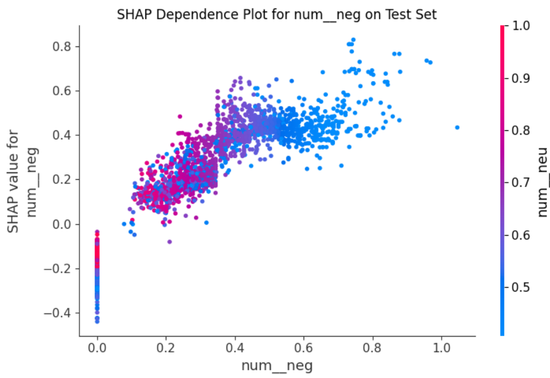

Fig. 14. SHAP Dependence Plot for Negative Sentiment on Test Set

The SHAP dependence plot for contracted word forms in Figure 15 demonstrates the relationship between the use of contractions (e.g., "won't," "can't") and the likelihood of a headline being classified as sensational. As the count of contracted word forms increases, the SHAP value also rises, suggesting that headlines with more contractions are more likely to be sensational. This trend becomes prominent around a contracted word count of 0.2-0.3, where the SHAP values show a sharp increase. The colour gradient representing the number of stop words further highlights that higher numbers of contracted forms and stop words both positively contribute to sensationalism detection.

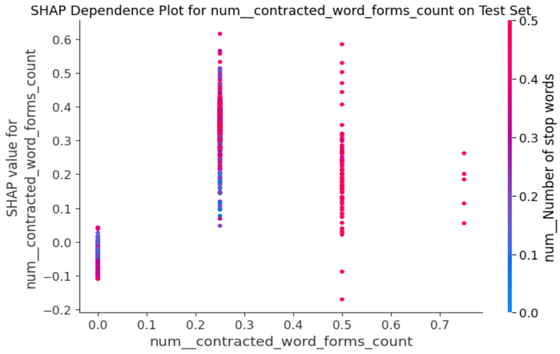

Fig. 15. SHAP Dependence Plot for Contracted Word Forms Count on Test Set

The SHAP dependence plot for double quotes count in Figure 16 illustrates the relationship between the use of double quotation marks and the model's classification of sensational headlines. As the count of double quotes increases, the SHAP value rises, indicating that headlines with more quoted text are more likely to be considered sensational. This relationship becomes evident around a double quote count of 0.3, where the SHAP values show a sharp increase. The interaction with stop words, shown by the colour gradient, further emphasizes the role of simple, quoted phrases in sensational headlines.

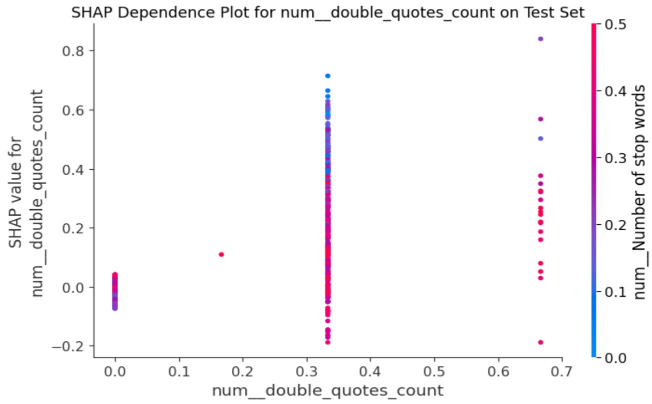

Fig. 16. SHAP Dependence Plot for Double Quotes Count on Test Set

In this analysis, we can observe the importance of various categories of linguistic features—lexical, syntactic, semantic, and readability—in detecting sensationalism. Lexical features, like the number of stop words and word count, emerged as key predictors of sensationalism, reflecting simpler language and shorter headlines' tendency to grab attention. Semantic features, such as sentiment (negative and compound scores), further highlight the emotional appeal present in sensational content. Additionally, syntactic features, like the use of contracted word forms and punctuation (e.g., double quotes), contribute significantly by pointing out informal or dramatic sentence structures. These diverse feature categories work together to provide a more comprehensive understanding of the linguistic strategies used in crafting sensational headlines.

Limitations

While this study has made significant strides in detecting sensationalism through linguistic feature analysis and the use of explainable machine learning models, several limitations must be acknowledged:

- **Bias in Data Sources**: The dataset used in this study primarily consists of headlines from well-known online news sources. This introduces potential bias, as these outlets may follow specific editorial standards that affect the tone, language, or level of sensationalism in their headlines. Smaller, regional, or niche media outlets, which may exhibit different sensationalism patterns, were not included, potentially limiting the generalizability of the results.

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index

- **Granularity of Annotation**: Sensation and emotion annotations rely heavily on linguistic markers and Likert scale ratings. While OpenAI's GPT-4 assisted in the annotation process, these models are trained on existing patterns, which may not capture more subtle variations in sensationalism. Additionally, human oversight, though incorporated, may introduce subjective bias in interpreting and categorizing sensational content, especially with borderline cases.
- **Simplification of Sensationalism**: The study simplifies sensationalism by focusing on lexical and syntactic features, leaving out cultural, contextual, and psychological factors that may influence the perception of sensationalism. Sensationalism is often context-dependent, and the same headline might be considered sensational in one context but not in another. The model does not account for these variations, limiting its effectiveness in understanding the full spectrum of sensationalism.
- **Language and Cultural Limitations**: The dataset is predominantly in English, which may limit the applicability of the findings to other languages or cultural contexts. Sensationalism can manifest differently across languages and cultures, with unique stylistic and rhetorical devices that are not captured in this study. Extending the dataset to include multilingual sources could improve the model's robustness in global news contexts.
- **Limitations of Sentiment and Readability Analysis**: While sentiment analysis and readability metrics (like Flesch-Kincaid) are useful indicators, they have their shortcomings. For instance, readability metrics may label emotionally charged but grammatically simple headlines as highly readable, without considering the emotional complexity of the content. Similarly, sentiment analysis tools like VADER may not accurately capture nuanced or sarcastic sentiments, potentially leading to misclassifications of sensational content.
- **Real-World Applicability**: Implementing this model in real-world applications, such as content moderation systems or journalistic ethics platforms, poses challenges. The dynamic nature of online news, where language trends and styles evolve rapidly, may require frequent retraining of the model to stay effective. Additionally, the system's effectiveness in contexts where multimedia elements (images, videos) accompany headlines has not been addressed.
- **Overreliance on Textual Features**: The study relies exclusively on textual features to detect sensationalism, which may overlook multimodal signals like images, videos, or layout, often used to amplify the sensational impact of news stories. In many cases, sensationalism is conveyed not only through text but through visuals and presentation style, which this model does not capture.
- **Computational Complexity**: The use of SHAP for model interpretability, while valuable for understanding feature contributions, introduces significant computational overhead. SHAP values, especially when calculated for large datasets, require substantial computational resources, making real-time implementation of sensationalism detection challenging, particularly in resource-constrained environments like mobile devices or real-time journalism platforms.
- **Lack of Human-Centred Evaluation**: While the study evaluates model performance through metrics such as precision, recall, and F1-score, it does not consider user studies or human-centred evaluations. Sensationalism detection systems would benefit from feedback from journalists, readers, and regulators to ensure the model aligns with human expectations and ethical standards in news reporting.
- **Feature Selection and Dimensionality**: Although steps were taken to optimize the feature set, the high-dimensional nature of textual data still presents challenges. Some potentially valuable features may have been overlooked or insufficiently explored due to the complexity of natural language and the necessity of dimensionality reduction.

Despite these limitations, the study provides a solid foundation for future research to address these challenges and further enhance the model's performance and applicability.

Mitigation Strategies
To address the identified limitations, several mitigation strategies can be implemented to enhance the model's robustness, applicability, and practicality. First, the issue of bias in data sources can be mitigated by expanding the dataset to include headlines from smaller, regional, and niche media outlets. This would provide a more comprehensive view of sensationalism patterns across diverse editorial contexts. Additionally, incorporating multilingual and culturally diverse headlines can address the language and cultural limitations, ensuring the model captures the stylistic and rhetorical variations of sensationalism across different regions.

For improving the granularity of annotation, advanced annotation techniques, such as semi-supervised learning and active learning, can be employed to refine the labelling process and reduce reliance on subjective human judgment. These methods, coupled with improved training datasets, can help capture subtle variations in sensationalism that may be missed in the current approach. Regular audits of annotated data with diverse reviewers can also help minimize subjective biases.
The simplification of sensationalism to lexical and syntactic features can be mitigated by integrating contextual, cultural, and psychological factors into the analysis. This could involve incorporating external knowledge sources, such as sentiment lexicons specific to different cultures or psychological studies on media impact, to build a more nuanced understanding of sensationalism. Incorporating domain adaptation techniques could also help tailor the model to context-specific interpretations.

To overcome limitations in sentiment and readability analysis, more advanced tools, such as fine-tuned BERT-based models for sentiment detection, can be utilized to better capture nuanced or sarcastic sentiments. Additionally, readability metrics could be enhanced to account for emotional complexity alongside grammatical simplicity, providing a more comprehensive analysis of sensational content. The challenge of real-world applicability can be addressed by designing dynamic models capable of periodic retraining to adapt to evolving language trends. Furthermore, incorporating multimodal analysis that considers images, videos, and layout alongside text would significantly enhance the detection of sensationalism. This approach would provide a holistic view of how sensationalism is conveyed in modern media.

To mitigate the computational complexity of SHAP, alternative interpretability techniques that balance efficiency and explainability, such as integrated gradients or LIME (Local Interpretable Model-Agnostic Explanations), can be explored. Developing lightweight versions of SHAP tailored for real-time applications could further reduce computational overhead and facilitate practical deployment in resource-constrained environments. Finally, introducing human-centred evaluations, such as user studies involving journalists, readers, and regulators, would ensure the model aligns with real-world expectations and ethical standards. Feedback from these stakeholders can guide refinements to the system, making it more practical and impactful. By implementing these mitigation strategies, the framework can be significantly improved to address its current limitations and broaden its applicability in detecting sensationalism across diverse media contexts.

CONCLUSION AND FUTURE WORK

This research has successfully developed a robust framework for detecting sensationalism in news headlines through the integration of linguistic feature analysis and advanced machine learning models. The creation of the novel SENS-HEAD dataset, featuring over 30,000 annotated headlines, provided a rich resource for exploring sensationalism. By leveraging features such as the number of stop words, sentiment polarity, readability, and syntactic patterns, we trained ensemble models, including XGBoost, CATBoost, and Random Forest, achieving balanced F1-scores across both sensational and non-sensational categories. The application of SHAP analysis enabled greater transparency, offering interpretability into how specific features influenced the model's predictions. The results demonstrated that certain features, such as the frequency of stop words and headline length, consistently played a critical role in predicting sensationalism. SHAP values highlighted that these features exhibited non-linear relationships with model output, indicating that shorter, simpler headlines with emotionally charged language were more likely to be classified as sensational. In terms of performance, CATBoost and XGBoost emerged as the most effective models, with F1-scores nearing 0.66 for both classes. The integration of superlative adjectives and an optimal threshold further refined model performance, illustrating the value of fine-tuning feature engineering in detecting exaggerated language commonly associated with sensationalism.

The policy impact of this research is significant, particularly in the realms of media regulation, ethical journalism, and content moderation. The insights derived from this study can inform policymakers and regulators on the prevalence and characteristics of sensationalism in news content, aiding the development of standards and guidelines to curb misleading or exaggerated reporting. The explainable nature of our framework also ensures alignment with ethical AI principles, allowing regulators and industry stakeholders to adopt transparent tools for moderating sensational content. By promoting accountability in media practices, this research contributes to fostering trust in news and mitigating the adverse societal impacts of sensationalism.

Given the computational demands and added complexity of contextual embeddings, our study prioritized creating a scalable and interpretable framework, with future work planned to explore these advanced methods. We also aim to expand the dataset to include multilingual and region-specific headlines, addressing cultural biases and enhancing model applicability. Incorporating multimodal data—such as images, videos, and user interaction metrics could further refine predictions by capturing visual and interactive cues tied to sensational content. Future efforts will focus on developing efficient real-time explainability techniques for practical deployment in settings like newsrooms. Additionally, investigating the psychological and societal impacts of sensationalism through longitudinal studies and designing personalized detection systems tailored to user preferences will promote ethical media consumption and AI transparency. Real-world evaluations and collaborations will be key to ensuring practical usability and ethical alignment in future advancements. By bridging the gap between research, practice, and policy, this study lays the groundwork for more responsible media practices and supports the broader goal of creating an informed and ethically aware media ecosystem.

Statements and Declarations

Conflict of Interest: The Authors have no competing interests related to the work submitted for publication.

Funding: No funding was received for this research.

Data Availability Statement: A dataset will be made available on request

Use of Generative AI: ChatGPT and Grammarly to assist with improving sentence ordering, reducing word count, and enhancing grammar. After using these tools, the authors meticulously reviewed and edited the content to ensure it met the required standards and take full responsibility for the final submission.

CRediT (Contributor Roles Taxonomy) Author Statement

Po-Hsuan Chang: Data curation, Methodology, Visualization, Investigation, Validation
Akshi Kumar: Conceptualization, Methodology, Writing- Original draft preparation., Supervision
Saurabh Raj Sangwan: Conceptualization, Writing-Reviewing and Editing

REFERENCES

1. ur Rahman, H. (2023). Media Ethics in the Era of Clickbait Journalism: Ethical Dilemmas and Solutions in Online Media. *Journal of Social Sciences Review*, *3*(4), 11-20.
2. Zuhroh, N. A., & Rakhmawati, N. A. (2020). Clickbait detection: A literature review of the methods used. *Register: Jurnal Ilmiah Teknologi Sistem Informasi*, *6*(1), 1-10.
3. Bazaco, Á., Redondo, M., & Sánchez-García, P. (2019). Clickbait as a strategy of viral journalism: conceptualisation and methods. *Revista Latina de Comunicación Social*, (74), 94.
4. Molek-Kozakowska, K. (2013). Towards a pragma-linguistic framework for the study of sensationalism in news headlines. *Discourse & Communication*, *7*(2), 173-197.
5. Wahl-Jorgensen, K. (2016). Emotion and journalism. *The SAGE handbook of digital journalism*, 128-143.
6. Messalas, A., Kanellopoulos, Y., & Makris, C. (2019, July). Model-agnostic interpretability with shapley values. In *2019 10th International Conference on Information, Intelligence, Systems and Applications (IISA)* (pp. 1-7). IEEE.
7. Kleemans, M., & Hendriks Vettehen, P. G. J. (2009). Sensationalism in television news: A review.
8. Biyani, P., Tsioutsiouliklis, K., & Blackmer, J. (2016). "8 amazing secrets for getting more clicks: Detecting clickbaits in news streams using article informality." Proceedings of the 30th AAAI Conference on Artificial Intelligence.
9. Brown, D. K., Harlow, S., García-Perdomo, V., & Salaverría, R. (2018). A new sensation? An international exploration of sensationalism and social media recommendations in online news publications. *Journalism*, *19*(11), 1497-1516.
10. Montejo, G. M., & Adriano, T. Q. (2018). A critical discourse analysis of headlines in online news portals. *Journal of Advances in Humanities and Social Sciences*, *4*(2), 70-83.
11. Xu, P., Wu, C. S., Madotto, A., & Fung, P. (2019). Clickbait? sensational headline generation with auto-tuned reinforcement learning. *arXiv preprint arXiv:1909.03582*.
12. Mourão, R. R., & Robertson, C. T. (2019). Fake news as discursive integration: An analysis of sites that publish false, misleading, hyperpartisan and sensational information. *Journalism studies*, *20*(14), 2077-2095.
13. Arbaoui, B., De Swert, K., & Van der Brug, W. (2020). Sensationalism in news coverage: A comparative study in 14 television systems. *Communication Research*, *47*(2), 299-320.
14. Naeem, B., Khan, A., Beg, M. O., & Mujtaba, H. (2020). A deep learning framework for clickbait detection on social area network using natural language cues. *Journal of Computational Social Science*, *3*(1), 231-243.
15. Allen, L. D., & Ayalon, L. (2021). “It’s pure panic”: The portrayal of residential care in American newspapers during COVID-19. *The Gerontologist*, *61*(1), 86-97.
16. Wasserman, H., Chuma, W., Bosch, T., Uzuegbunam, C. E., & Flynn, R. (2021). South African newspaper coverage of COVID-19: A content analysis. *Journal of African Media Studies*, *13*(3), 333-350.
17. Ottwell, R., Puckett, M., Rogers, T., Nicks, S., & Vassar, M. (2021). Sensational media reporting is common when describing COVID-19 therapies, detection methods, and vaccines. *Journal of Investigative Medicine*, *69*(6), 1256-1257.
18. Liepukhova, N. I., & Shcherbak, O. M. (2021). Lexico-semantic means of sensationalism in headings of German-language media texts.
19. González Esparza, H. J., Florencia, R., Díaz Román, J. D., & Mendoza-Carreón, A. (2023). Text Classifier of Sensationalist Headlines in Spanish Using BERT-Based Models. In *Innovations in Machine and Deep Learning: Case Studies and Applications* (pp. 109-131). Cham: Springer Nature Switzerland.

Print ISSN: 2517-276X
Online ISSN: 2517-2778
https://bjmas.org/index.php/bjmas/index
20. Hamby, A., Kim, H., & Spezzano, F. (2024). Sensational stories: The role of narrative characteristics in distinguishing real and fake news and predicting their spread. *Journal of Business Research*, *170*, 114289.
21. Khawar, S., & Boukes, M. (2024). Analyzing Sensationalism in News on Twitter (X): Clickbait Journalism by Legacy vs. Online-Native Outlets and the Consequences for User Engagement. *Digital Journalism*, 1-21.
22. Zhang, S., & Kejriwal, M. (2019, August). Concept drift in bias and sensationalism detection: An experimental study. In *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining* (pp. 601-604).
23. Joshi, A., Kale, S., Chandel, S., & Pal, D. K. (2015). Likert scale: Explored and explained. *British journal of applied science & technology*, *7*(4), 396-403.
24. Tong, Z., & Zhang, H. (2016, May). A text mining research based on LDA topic modelling. In *International conference on computer science, engineering and information technology* (pp. 201-210).
25. Heylighen, F., & Dewaele, J. M. (1999). Formality of language: definition, measurement and behavioral determinants. *Interner Bericht, Center "Leo Apostel", Vrije Universiteit Brüssel*, *4*(1).
26. Flesch, R. (2007). Flesch-Kincaid readability test. *Retrieved October*, *26*(3), 2007.
27. Aizawa, A. (2003). An information-theoretic perspective of tf–idf measures. *Information Processing & Management*, *39*(1), 45-65.
28. Shlens, J. (2014). A tutorial on principal component analysis. *arXiv preprint arXiv:1404.1100*.
29. Kumar, A., Jain, D. K., Mallik, A., & Kumar, S. (2024). Modified node2vec and attention based fusion framework for next POI recommendation. *Information Fusion*, *101*, 101998.
30. Kumar, A., Mallik, A., & Kumar, S. (2024). HumourHindiNet: Humour detection in Hindi web series using word embedding and convolutional neural network. *ACM Transactions on Asian and Low-Resource Language Information Processing*, *23*(7), 1-21.
31. González Esparza, H. J., Florencia, R., Díaz Román, J. D., & Mendoza-Carreón, A. (2023). Text Classifier of Sensationalist Headlines in Spanish Using BERT-Based Models. In *Innovations in Machine and Deep Learning: Case Studies and Applications* (pp. 109-131). Cham: Springer Nature Switzerland.
32. Linardatos, P., Papastefanopoulos, V., & Kotsiantis, S. (2020). Explainable ai: A review of machine learning interpretability methods. *Entropy*, *23*(1), 18.