

# Explaining Humour Style Classifications: An XAI Approach to Understanding Computational Humour Analysis

Mary Ogbuka Kenneth<sup>1</sup>, Foad Khosmood<sup>2</sup>, Abbas Edalat<sup>1</sup>

<sup>1</sup>Algorithmic Human Development group, Department of Computing, Imperial College London, UK

<sup>2</sup>Computer Engineering Department, California Polytechnic State University, USA

Corresponding author: Mary Ogbuka Kenneth, [m.kenneth22@imperial.ac.uk](mailto:m.kenneth22@imperial.ac.uk)

## Abstract

Humour styles can have either a negative or a positive impact on well-being. Given the importance of these styles to mental health, significant research has been conducted on their automatic identification. However, the automated machine learning models used for this purpose are black boxes, making their prediction decisions opaque. Clarity and transparency are vital in the field of mental health. This paper presents an explainable AI (XAI) framework for understanding humour style classification, building upon previous work in computational humour analysis. Using the best-performing single model (ALI+XGBoost) from prior research, we apply comprehensive XAI techniques to analyse how linguistic, emotional, and semantic features contribute to humour style classification decisions. Our analysis reveals distinct patterns in how different humour styles are characterised and misclassified, with particular emphasis on the challenges in distinguishing affiliative humour from other styles. Through detailed examination of feature importance, error patterns, and misclassification cases, we identify key factors influencing model decisions, including emotional ambiguity, context misinterpretation, and target identification. The framework demonstrates significant utility in understanding model behaviour, achieving interpretable insights into the complex interplay of features that define different humour styles. Our findings contribute to both the theoretical understanding of computational humour analysis and practical applications in mental health, content moderation, and digital humanities research.

## Keywords

Explainable AI; Computational Humour; Humour Style Classification; Natural Language Processing; Machine Learning Interpretability; Digital Humanities

## I INTRODUCTION

Humour is the tendency to experience or provoke laughter or provide amusement through written or spoken words [Sen, 2012]. It plays a vital role in interpersonal interactions, emotional expression [Amjad and Dasti, 2022], and psychological well-being [Martin et al., 2003, Chen and Martin, 2007, Edalat, 2023, Martin and Ford, 2018]. Different humour styles—self-enhancing, self-deprecating, affiliative, and aggressive—carry distinct emotional undertones and social implications [Martin et al., 2003, Kuiper et al., 2016]. While affiliative humour fosters positive social interactions, aggressive humour may strain relationships through its potential to offend or demean [Anderson and Di Tunnariello, 2016]. Understanding these styles has significant implications across multiple domains, including mental health [Edalat et al., 2024, Martin and Ford, 2018, Martin et al., 2003], content moderation [Zhu et al., 2022, Sari, 2016], and artificial intelligence (AI) [Kenneth et al., 2024, Kenneth Ogbuka et al., 2024]. However, computational hu-

humour style recognition presents unique challenges in natural language processing (NLP) due to its subjective nature and complex psychological underpinnings [Kenneth et al., 2024, Kazienko et al., 2023, Amjad and Dasti, 2022]. Even though recent developments have been effective in categorising various humour styles, these models' decision-making procedures remain mainly unknown. Understanding how and why machine learning models categorise various humour styles is crucial for advancing both computational linguistics and digital humanities research [Cortinas-Lorenzo and Lacey, 2024].

Recent work by Kenneth Ogbuka et al. [2024] established baseline performance for humour style classification across four categories: self-enhancing, self-deprecating, affiliative, and aggressive humour. Their two-model approach achieved promising results, particularly with the General Text Embeddings Upgraded (ALI) + XGBoost for the single model configuration and the Multilingual E5 Text Embeddings (MUL) + XGBoost + ALI+XGBoost for the two-model configuration. However, the interpretability of these classifications—understanding which features and patterns drive the model's decisions—remains an open challenge.

Explainable AI (XAI) has emerged as a solution to address the opacity of traditional ML models by providing insights into their decision-making processes [Lyu et al., 2024]. Through techniques such as Local Interpretable Model-Agnostic Explanations (LIME) [Ribeiro et al., 2016] and Shapley Additive Explanations (SHAP) [Lundberg and Lee, 2017], XAI enables researchers to identify key features driving predictions. This interpretability is especially crucial for recognising humour styles, as the interaction of linguistic patterns, emotional tones, and semantic nuances necessitates a thorough comprehension. For instance, identifying features such as sarcasm, sentiment contrasts, or emotions can provide more detailed information about why a certain text is categorised as affiliative or self-deprecating.

This paper extends the work of Kenneth Ogbuka et al. [2024] by introducing an XAI framework for humour style classification. By applying XAI to the best-performing single model (ALI+XGBoost) reported in their research, we provide detailed interpretability that highlights the influence of linguistic and emotional features on model predictions. This approach not only makes the classification of humour styles more transparent, but it also gives researchers practical insights that allow them to further investigate the role of humour in communication and psychological well-being.

The primary contributions of this paper are as follows:

1. Development of a comprehensive XAI framework tailored to humour style recognition.
2. Detailed analyses of how linguistic and affective features contribute to humour style classification.
3. Insights into the practical implications of these findings for researchers studying humour and its applications.

By addressing the critical need for explainability, this study bridges the gap between computational performance and interpretability in the field of humour style recognition.

## II RELATED WORKS

The development of explainable approaches to computational humour analysis converges three research streams: general XAI methodologies for text classification, interpretable humour analysis models, and explainable style classification approaches. This section examines these areas' contributions to computational humour recognition interpretability, progressing from foundational XAI methods to specific style-based classifications.

## 2.1 General XAI Methods for Text Classification

Recent advances in XAI for text classification demonstrate varied approaches to model interpretability. Pérez-Landa et al. [2021] combined emotional, sentiment, syntactic, and lexical features for xenophobic content detection, achieving F1-scores of 76.8% and 73.4% on different datasets. Although their keyword-based pattern matching may face generalisation challenges, their success with emotion and sentiment features informs our approach to humour style explanation.

Chowdhury et al. [2021] applied LIME to explain black-box sentiment analysis models, achieving 72% accuracy with Long Short-Term Memory (LSTM) networks and FastText embeddings. Their implementation revealed word-level contributions to sentiment predictions through probability scores and contextual relationships, demonstrating LIME's effectiveness for deep learning interpretability.

Ahmed et al. [2022] integrated attention mechanisms with fuzzy logic rules for interpretable sentiment analysis, achieving 89% F1-score. Their dual-layer approach combined local explanations with global interpretability through attention weights and fuzzy rules. While effective for mental health applications, questions remain about adapting such hybrid systems for subjective tasks like humour classification, where decision boundaries may be less clearly defined.

In humour-adjacent tasks, Ortega-Bueno et al. [2022] employed attention mechanisms for irony detection in Spanish, demonstrating that different attention types focus on distinct linguistic features. However, their approach prioritised performance over explainability, underscoring the need for XAI methods that consider both linguistic feature utilisation and cross-cultural variations.

## 2.2 Interpretable Models for Humour Analysis

Early interpretable humour analysis focused on transparent feature development. Zhang et al. [2017] proposed interpretable features for humour recognition: contextual knowledge (modelling semantic relationships), affective polarity (quantifying emotional impact), and subjectivity (capturing personal interpretation). Their CASHCF (Combined All Semantic and Human-Centred Features) model improved performance while maintaining transparency but highlighted challenges in capturing deep semantic relationships in an explainable manner.

Mann and Mikulandric [2024] addressed pun detection, humour classification (irony, exaggeration, incongruity, self-deprecating, wit), and English-French joke translation using TF-IDF features and MarianMT. They found wit the easiest to classify, while irony and exaggeration were more challenging, supporting our style-specific XAI approach. However, their reliance on basic features limits interpretability. Our framework advances this by incorporating semantic, emotional, and linguistic analyses.

De Marez et al. [2024] introduced THInC (Theory-driven Humour Interpretation and Classification), a framework grounding humour detection in psychological theories. Using GA2M classifiers aligned with humour theories (superiority, relief, incongruity, and incongruity resolution), they achieved an F1 score of 85% through theory-informed proxy features like emotional bursts and sentiment shifts. Their mapping of computational features to theoretical constructs informs our XAI approach, though their focus is on binary humour classification rather than style differentiation.

Mahajan and Zaveri [2024] proposed a humour detection model using ensemble learning and Berger’s humour typology features (e.g., emotive, incongruity, intensity). Their stacking-based ensemble achieved 85.68% accuracy and a 72.57% F1-score on Yelp reviews, with ablation studies highlighting the value of theory-driven features. While focused on binary detection, their use of typology offers insights for interpretable humour analysis.

Chen et al. [2024] introduced TalkFunny, a Chinese humour response dataset with ”chain-of-humour” annotations and mind maps explaining response generation. Their PLM-based framework, integrating interpretability tools, showed improved performance. While focused on humour generation, their approach supports the value of explainable components, aligning with our goals for interpretability in humour AI.

### 2.3 Explainable Approaches to Style Classification

Style classification has evolved from binary to multi-class approaches. Abulaish and Kamal [2018] developed a two-layer system for self-deprecating sarcasm detection (94% F1-score), using explicit linguistic rules and feature-based classification. Despite its success in binary classification, their work underscored the need for more advanced explainability to handle multiple humour styles. Kamal and Abulaish [2020] expanded this work with three explicit feature categories (self-deprecating patterns, exaggeration markers, and word embeddings) for self-deprecating humour detection, achieving F1-scores of 62%–87% across datasets.

Kenneth Ogbuka et al. [2024] introduced a dataset and a two-model approach for humour style classification, targeting four styles (self-deprecating, self-enhancing, affiliative, aggressive) and non-humour. They addressed challenges in distinguishing affiliative and aggressive humour by using a four-class classification followed by binary discrimination, achieving a 78.6% F1-score with improved differentiation. However, their work left unanswered questions about misclassification reasons between styles—a gap our XAI framework seeks to fill.

While significant progress has been made in humour classification and XAI techniques, two significant gaps persist: most interpretable approaches focus on binary rather than style-specific classification, and current style-specific approaches lack comprehensive explanatory frameworks for linguistic and emotional feature interaction. Our framework addresses these gaps through a unified approach to humour style classification explanation, supporting both theoretical understanding and practical applications. A comprehensive summary of related works and their contributions is presented in Table 5 (Appendix A).

## III METHODOLOGY

This study develops an XAI framework to analyse humour style classification by examining model predictions, linguistic features, and emotional patterns. Our methodology comprises three main components: dataset and model selection, prediction analysis using LIME, and comprehensive feature analysis across linguistic, affective, and contrast patterns.

### 3.1 Dataset and Classification Model

This study utilises the dataset and ALI+XGBoost model from Kenneth Ogbuka et al. [2024], which achieved 77.8% accuracy and 77.3% F1-score. The dataset comprises 1,463 instances gathered from multiple sources:

- 983 jokes from different websites (Reader’s Digest, Parade, Bored Panda, Laugh Factory, Pun Me, Independent, Cracked, Reddit, Tastefully Offensive, and BuzzFeed), labelled based on original website tags and humour theory.

- 280 non-humorous text instances from the ColBERT dataset [Annamoradnejad and Zoghi, 2020]
- 200 instances from the Short Text Corpus <sup>1</sup>, including 150 jokes and 50 non-jokes

To ensure annotation quality and mitigate potential biases, six Ph.D. candidates from Africa, Asia, and Europe independently annotated 200 randomly selected instances from the Short Text Corpus. Each set of 100 samples was reviewed by three annotators, with final labels determined by majority vote. The dataset contains five categories: self-enhancing (298 instances), self-deprecating (265), affiliative (250), aggressive (318), and neutral humour (332). Text lengths vary considerably, ranging from 4 to 229 words, with a total vocabulary of 4,506 unique words. Each humour style has the following linguistic characteristics:

- **Self-enhancing:** 1,181 unique words, with lengths ranging from 5 to 35 words, median of 13 and mode of 10 words
- **Self-deprecating:** 1,365 unique words, with lengths ranging from 4 to 148 words, median of 15 and mode of 12 words
- **Affiliative:** 1,376 unique words, with lengths ranging from 5 to 229 words, median of 15 and mode of 13 words
- **Aggressive:** 1,449 unique words, with lengths ranging from 4 to 112 words, median of 15 and mode of 12 words
- **Neutral:** 1,856 unique words, with lengths ranging from 4 to 34 words, median of 10 and mode of 9 words

While the length distributions show positive skewness with some long examples (particularly in affiliative humour), the central tendencies are quite similar across humour styles, with medians ranging from 10-15 words and modes from 9-13 words. The neutral category shows the most compact distribution (4-34 words), while humorous categories occasionally include longer examples. However, since the majority of instances cluster around similar lengths, these outliers are unlikely to impact classification performance significantly.

To illustrate these categories, representative examples from the humour styles dataset are presented in Figure 1:

Using this dataset, Kenneth Ogbuka et al. [2024] developed two classification approaches: a single-model and a two-model system. The single-model directly classified input text into one of the five categories (self-enhancing, self-deprecating, affiliative, aggressive, and neutral). Due to challenges in distinguishing affiliative from aggressive humour, they developed a two-model approach that:

1. First classifies text into four categories (self-enhancing, self-deprecating, neutral, and a combined affiliative/aggressive class).
2. Then uses a separate binary classifier to distinguish between affiliative and aggressive instances for texts classified in the combined class.

In this study, we analyse the prediction of the best-performing single-model (ALI+XGBoost) to understand these classification challenges from the perspective of explainable AI.

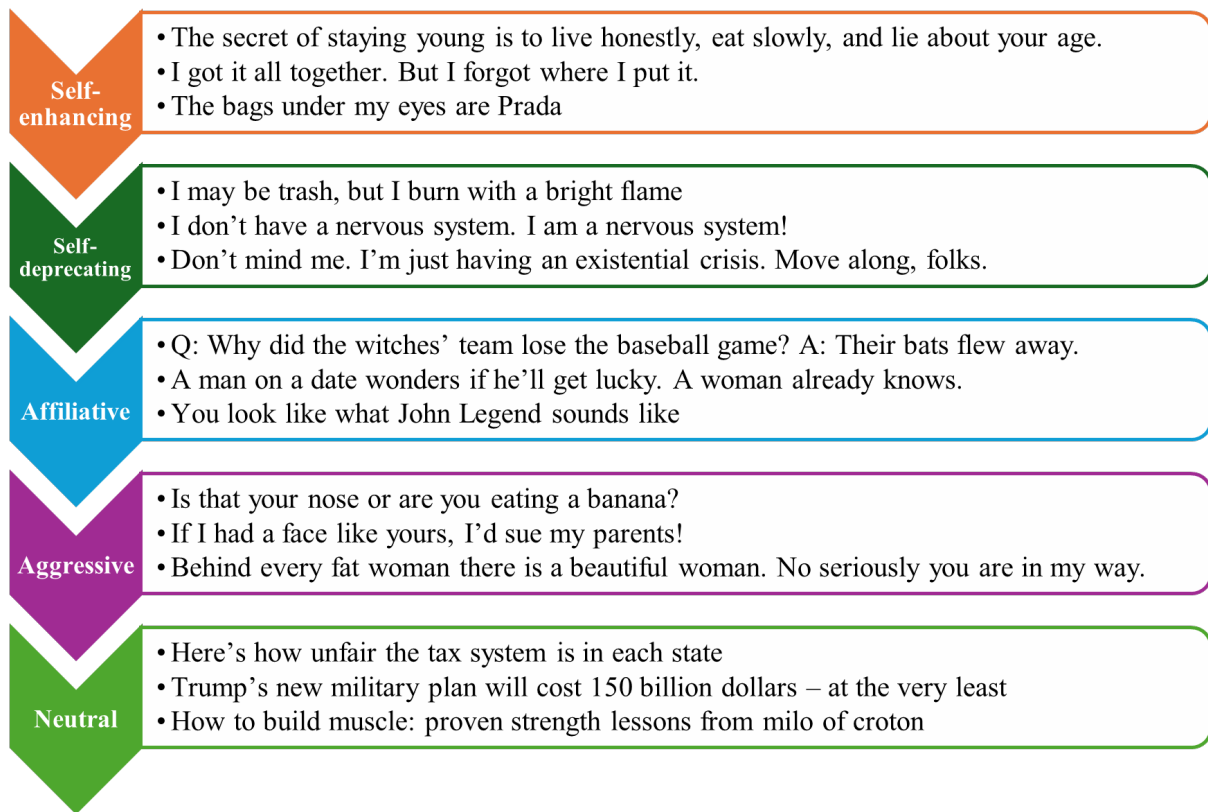
### 3.2 Model Prediction Analysis

We employ LIME (Local Interpretable Model-agnostic Explanations) due to its advantages for word embedding-based classification [Ribeiro et al., 2016]. LIME provides word-level interpretability through local explanations while maintaining model-agnostic analysis and computational efficiency. Our analysis progresses through:

<sup>1</sup>Short Text Corpus (<https://github.com/CrowdTruth/Short-Text-Corpus-For-Humor-Detection>)



Figure 1: Examples for Each Category in the Dataset



1. Individual prediction explanation using LIME
2. Feature importance extraction from XGBoost
3. LIME result visualization
4. Analysis of confidence score across humour styles

### 3.3 Feature Analysis

We categorise the patterns extracted to explain the classification model into three groups: linguistic patterns, affective patterns, and contrast patterns. Our analysis leverages multiple feature detection algorithms, each with different capabilities and limitations.

#### 3.3.1 Linguistic Patterns

Linguistic patterns capture the language elements that contribute to humour's comedic effect. These patterns manipulate structure, context, and style to create humorous situations [Kenneth et al., 2024]. The key linguistic patterns analysed include:

**Sound Patterns** . Sound patterns exploit auditory features of language, such as rhyme, alliteration, and homophony.

- **Rhyme:** The repetition of similar sounds at the ends of words. Examples include *sight* and *flight*, *sad* and *mad*, *cat* and *hat*. In our analysis, we used the “pronouncing” library with the CMU Pronouncing Dictionary to detect rhymes. Although the dictionary effectively covers standard English pronunciations, it may not capture slang or neologisms present in the dataset. For instance, out of the 4,506 unique words in the dataset, 90.75% (4,089 words) were found in the CMU Pronouncing Dictionary. Unmatched examples include: ‘*launchalot*’, ‘*aronofskys*’, ‘*500000*’, ‘*ahappy*’, ‘*admitted*’, ‘*kanye*’, ‘*houseplant*’, ‘*idk*’, ‘*behaviours*’

- **Alliteration:** The repetition of initial consonant sounds in closely placed words, such as in "*Crazy cats create chaos*". Implemented using NLTK's phonetic detection and spaCy's tokenisation, the algorithm identifies words sharing initial phonemes via the CMU Pronouncing Dictionary. However, this approach has limitations:
  - It considers only the first phoneme, occasionally grouping words incorrectly (e.g., "*when*" and "*one*").
  - It counts repeated words as alliteration.
  - It ignores word proximity, leading to false positives.
- **Homophones:** Words that sound identical but have different meanings or spellings. Examples include: *cell* and *sell*, *sea* and *see*, *to*, *too*, and *two*. Detection was implemented using WordHoard's homophone identification combined with WordNet synsets for semantic differentiation. While this approach performed well, it remains limited by WordNet's vocabulary size.

**Wordplay.** This category analyses deliberate manipulations of word meanings to create humour.

- **Puns:** Words with similar pronunciations but different meanings, identified through phonetic similarity and semantic difference analysis. For instance, "*A bicycle can't stand on its own because it is **two-tired**.*" Here, the bolded word is a homophonic pun, replacing too-tired with two-tired. Our implementation detects homophonic puns using phonetic matching (using 'pronouncing') and meaning differentiation (using WordNet). Current limitations include:
  - Only detects single-word homophonic puns with exact sound matches
  - Cannot identify near-homophones, multi-word puns, or contextual wordplay
  - Misses puns based on multiple meanings of the same word
- **Synsets:** Words or phrases that can be interpreted in multiple ways, detected using WordNet synset analysis. Each word in the dataset is matched with its WordNet synsets, extracting all possible meanings. Approximately 88.50% (3,988 of 4,506 words) of the vocabulary had synset coverage. Words like '*ourselves*', '*launchalot*', '*aronofskys*', '*since*', and '*500000*' were not included in WordNet.
- **Syllabic Structure:** Patterns in syllable counts and distributions within the text, analysed for their role in creating rhythm or emphasis. This was implemented using the Pyphen dictionary's hyphenation patterns as a proxy for syllable counting. In this approach, the algorithm counts hyphenation points in each word using Pyphen's default hyphenation rules. This method has inherent limitations, as hyphenation patterns do not always correspond directly to syllabic boundaries. Potential limitations include:
  - Reliance on hyphenation rules rather than true syllabic structure
  - No special handling for compound words or complex morphology
  - Limited accuracy for abbreviations, numerals, and non-standard English words
  - Dependency on Pyphen's dictionary coverage

**Structural Elements.** Structural elements focus on how linguistic arrangements contribute to humour.

- **Self-references:** Details with the identification of first-person pronouns and personal narratives. We used basic tokenisation to identify explicit self-references ('*i*', '*me*', '*my*', '*mine*', '*myself*') and associated context. However, implicit self-references may be missed.
- **Parts of speech (POS):** Analysis of the distribution and arrangement of word types, such as nouns, verbs, and adjectives. POS tagging was implemented using TextBlob, which offers basic functionality but lacks handling for informal language or social media text.

- **Grammar Complexity:** Measurement of sentence complexity by analysing the frequency of specific syntactic dependencies, such as clausal complements (ccomp), open clausal complements (xcomp), and adverbial clause modifiers (advcl).

These linguistic patterns were systematically extracted using NLTK, WordNet, [wordhoard](#), and custom phonetic analysis algorithms. By detecting these patterns, the analysis explains how different linguistic devices contribute to humour styles, providing interpretable features for the classification model.

### 3.3.2 Affective Patterns

Affective patterns relate to feelings, moods, and attitudes, as defined by the *Collins English Dictionary*. In humour style analysis, these patterns help distinguish whether a statement is humorous, sarcastic, or offensive [Kenneth et al., 2024]. Key affective patterns analysed include:

**Sarcasm.** Sarcasm detection was performed using a [RoBERTa-based](#) model trained for sarcasm identification, achieving 60.7% accuracy. The model provides binary classification (sarcastic/non-sarcastic) and probability scores, capturing subtle verbal irony often found in specific humour styles.

**Sentiment.** Sentiment analysis employed a multi-layered approach using [RoBERTa](#) (93.2% average accuracy) [Hartmann et al., 2023] and TextBlob models to assess:

- Dominant sentiment (positive, negative, neutral) with confidence scores
- Sentiment strength measured through positive-negative score differential
- Polarity (-1 to 1) indicating sentiment direction and intensity
- Subjectivity and objectivity metrics to gauge emotional content

**Emotion.** Emotional content was analysed using a [DistilBERT-based](#) model that classifies text into six discrete emotion categories (joy, anger, sadness, fear, love, and surprise) with an accuracy of 93.8% and F1-score of 93.79%. This fine-grained emotional analysis reveals distinctive affective patterns across different humour styles.

### 3.3.3 Contrast Patterns

Our analysis of contrast patterns focuses on semantic and sentiment-based contradictions that generate humorous effects through the following components:

**Sentiment Contrasts.** We examined opposing emotional valences at two distinct levels:

- **Sentence-level contrasts:** Sequential sentences were analysed for opposing sentiment polarities, capturing dramatic shifts in emotional tone.
- **Word-level contrasts:** Content words (adjectives, adverbs, verbs, and nouns) were evaluated using SentiWordNet to identify internal sentiment conflicts within sentences.

**Semantic Elements.** We investigated semantic relationships and conflicts through:

- **Exaggeration markers:** Systematic identification of absolute terms and extreme descriptors that amplify narrative elements.
- **Intensification patterns:** Analysis of linguistic intensifiers that heighten semantic impact.
- **Semantic incongruity:** Quantification of conceptual contradictions using WordNet similarity metrics to identify semantically distant word pairs.

Having established our methodological framework and feature detection capabilities, we now turn to the analysis of our results and their implications for humour style classification.



## IV RESULTS AND DISCUSSION

This section presents an analysis of the ALI+XGBoost model’s performance in humour style classification, followed by a detailed examination of linguistic mechanisms and error patterns. We first evaluate the model’s classification metrics and confidence scores across different humour styles, then analyse the distinctive linguistic features characterising each style, and finally examine misclassification patterns to understand the model’s limitations. The complete LIME visualisations, linguistic analysis, and affective analysis results for all 293 examined instances are available in our public repository <sup>2</sup>.

### 4.1 Model Performance Analysis

The performance analysis of the ALI+XGBoost model highlights distinct patterns in classification accuracy and confidence across different humour styles. The model achieved an overall accuracy of 78% and a macro-average F1-score of 77%, with style-specific variations:

- Self-enhancing humour exhibited the highest precision (0.82) and F1-score (0.83).
- Neutral content attained the highest recall (0.93) and a strong F1-score (0.85).
- Affiliative humour had the lowest recall (0.58), underscoring challenges in identifying this style.
- Aggressive humour showed balanced precision (0.75) and recall (0.74).

Confidence scores also varied significantly across styles ( $p < 0.0001$ ): Self-enhancing humour had the highest average confidence (0.889), followed by neutral (0.887), self-deprecating (0.811), aggressive (0.780), and affiliative humour (0.748).

### 4.2 Linguistic Mechanism Analysis

Analysis of linguistic mechanisms revealed distinct patterns across humour styles, as summarised in Table 1. We examined four key dimensions: syllabic complexity, semantic conflicts, homonym usage, and exaggeration patterns.

#### 4.2.1 Syllable Complexity

Neutral content exhibited significantly higher syllabic complexity (mean: 1.355, SD: 0.245), suggesting more formal or complex language use. In contrast, aggressive humour demonstrated the lowest complexity (mean: 1.148, SD: 0.118), indicating a preference for simpler, more direct language. Other humour styles showed intermediate complexity levels, with self-enhancing (mean: 1.206) and self-deprecating (mean: 1.192) humour showing similar patterns.

#### 4.2.2 Semantic and Structural Elements

Affiliative humour exhibited the highest frequency of semantic conflicts (mean: 29.755, SD: 114.755), significantly surpassing other styles. This suggests frequent use of wordplay and unexpected combinations. Self-deprecating humour followed with a mean of 18.717 (SD: 47.830). In contrast, neutral and self-enhancing content demonstrated markedly lower frequencies of semantic conflicts (mean: 7.025, SD: 7.347, and mean: 7.557, SD: 7.274, respectively).

Homonym usage patterns revealed significant differences between humorous and neutral content. Social forms of humour—affiliative and aggressive—showed the highest homonym frequencies (mean = 6.796, SD = 4.743 and mean = 6.667, SD = 3.772, respectively), while neutral content demonstrated significantly lower usage (mean = 2.662, SD = 1.987). This indicates that wordplay involving multiple word meanings is a common feature across humour styles.

---

<sup>2</sup>XAI humour styles : [https://github.com/MaryKenneth/XAI\\_humour\\_styles/tree/main](https://github.com/MaryKenneth/XAI_humour_styles/tree/main)

Predicted Class	Syllable Complexity				Semantic Conflict Count				Homonym Count				Exaggeration Count			
	Mean	Std	Min	Max	Mean	Std	Min	Max	Mean	Std	Min	Max	Mean	Std	Min	Max
Affiliative	1.183	0.118	1.0	1.48	29.755	114.755	0	794	6.796	4.743	1	33	1.388	1.304	0	5
Aggressive	1.148	0.118	1.0	1.50	13.439	37.424	0	231	6.667	3.772	2	22	1.456	1.593	0	9
Neutral	1.355	0.245	1.0	2.00	7.025	7.347	0	48	2.662	1.987	0	11	0.738	0.853	0	3
Self-deprecating	1.192	0.171	1.0	2.00	18.717	47.830	0	294	6.087	3.817	0	22	1.848	1.738	0	8
Self-enhancing	1.206	0.143	1.0	1.64	7.557	7.274	0	30	4.902	2.350	0	12	0.885	0.950	0	4

Table 1: Complexity Statistics

### 4.2.3 Rhetorical Devices

Exaggeration patterns varied significantly across styles. Self-deprecating humour exhibited the highest frequency of exaggeration (mean = 1.848, SD = 1.738), followed by aggressive humour (mean = 1.456, SD = 1.593). Neutral and self-enhancing content showed significantly lower exaggeration frequencies (mean = 0.738, SD = 0.853 and mean = 0.885, SD = 0.950, respectively).

These findings suggest that humour styles employ distinctive combinations of linguistic mechanisms. Social forms of humour (affiliative and aggressive) predominantly utilise wordplay devices, while self-directed styles demonstrate varied complexity patterns. Neutral content maintains higher linguistic complexity but shows reduced usage of rhetorical devices.

### 4.2.4 Mechanism Correlations and Interactions

Analysis of correlations between linguistic mechanisms revealed complex interaction patterns in humour construction. We employed both Pearson and Spearman correlation analyses as shown in Figure 2 to understand these relationships more thoroughly, as Spearman correlations are more robust to outliers and can detect non-linear relationships.

#### Correlation Disparity and Outlier Effects

The most striking finding emerged in the disparity between Pearson and Spearman correlations, particularly for semantic conflict relationships. While Pearson correlation suggested a strong linear relationship between semantic conflicts and rhyme ( $r = 0.95$ ), the much lower Spearman correlation ( $r = 0.32$ ) revealed this relationship to be heavily influenced by outliers. This disparity manifests across different scaling levels:

- **High-density outliers:** 794 conflicts/181 rhymes; 294 conflicts/44 rhymes; 231 conflicts/27 rhymes
- **Mid-range instances:** 48 conflicts/7 rhymes; 30 conflicts/5 rhymes
- **Typical cases:** Majority fall below 30 conflicts; Variable rhyme counts (0-5)

This distribution pattern is illustrated by two contrasting joke examples:

**A typical case example:** *“Reminder that Winnie the Pooh wore a crop top with no pants and ate his fave food and loved himself. So you can too.”*

- 4 rhyming pairs: (“Pooh/you”, “Pooh/too”, “crop/top”, “you/too”)
- 5 semantic conflict: memory aid vs physical action (“reminder/wore”), memory aid vs emotional satisfaction (“reminder/loved”), clothing vs eating (“wore/food”), physical vs emotional satisfaction (“food/loved”), and the phrase “So you can too” suggests an invitation, but applying it to “food” and “loved himself” creates ambiguity (“food/too”).

**High rhyme example:** *“This is your mother. I just texted you but I don’t know how to make the facey-things so... happy face at the end.”*

- 4 distinct rhyming pairs (“you/do”, “you/to”, “do/to”, “know/so”)
- No significant semantic conflicts

Comparison of Correlation Methods for Humor Mechanisms

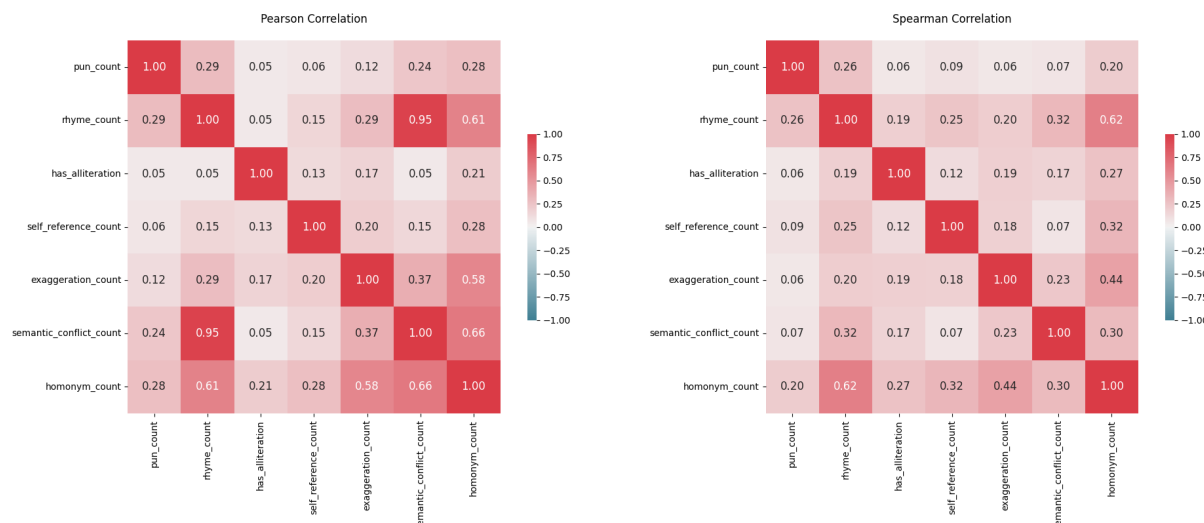


Figure 2: Mechanism Pearson and Spearman Correlations

### Stable Correlations and Robust Relationships

While many correlations weakened under Spearman analysis, some relationships demonstrated stability across both methods. Most importantly, the relationship between rhyme and homonym remained consistent (Pearson: 0.61, Spearman: 0.62), suggesting a robust underlying connection independent of outliers. Homonyms also showed stable correlations with exaggeration (Pearson: 0.58, Spearman: 0.44) and semantic conflicts (Pearson: 0.66, Spearman: 0.30).

Consider this example that demonstrates stable mechanism interactions: *“I forgot you can’t make depression jokes outside of twitter Imao my coworker was like ‘you ready for this year to be over?’ I was like I’m ready for this life to be over’ he was like bro what”*. This joke exhibits stable correlations between multiple features:

- 7 rhyming pairs such as “I/my”, “you/to”, “be/he”
- 6 homonyms like “I”: [“eye”, “aye”, “ai”]; “you”: [“yew”, “ewe”]
- 10 semantic conflicts supporting wordplay such as ‘forgot/depression’, “forgot/year”, “forgot/life”, “make/year”
- 3 self-references integrated with sound patterns- “I/my/I”
- 12 alliteration instances like “F”: [“forgot”, “for”]; “Y”: [“you”, “year”]; “M”: [“make”, “my”]; “T”: [“Twitter”, “to”]

### Independent Mechanisms

Several mechanisms demonstrated operational independence. Self-referential content maintained weak to moderate correlations with other features –Pearson range:  $(0.06 \leq r \leq 0.28)$  and similar stability range for Spearman, suggesting it functions largely autonomously in humour construction. Similarly, alliteration showed minimal correlation with other mechanisms (Pearson correlations  $<0.21$  and Spearman correlations  $<0.27$  across all comparisons), indicating its role as an independent supplementary device rather than a core component of humour construction.

Example joke: *“My manager asked if I take constructive criticism and I said yes while wiping away my teary eyes.”*

In the above joke, self-reference is present in four instances: “my/manager”, “I/take”, “I/said”,

and “my/teary”. These create humour through a personal narrative, establishing a relatable and emotional context without requiring support from other mechanisms.

Alliteration appears in six groups, such as the “M” sounds in [“my,” “manager,” “my”] and the “T” sounds in [“take,” “teary”]. While alliteration enhances the rhythmic flow of the joke, it functions as a supplementary device, adding aesthetic value without depending on other features for its effect.

The semantic structure of the joke contains 25 semantic conflicts such as “manager/criticism”, “manager/teary”, and “manager/yes”, which operate independently of sound patterns. Looking at these three example conflicts: 1) the manager’s formal criticism meeting an emotional response - demonstrating the professional-personal divide that operates independently of phonetic features; 2) the stark contrast between professional authority and teary vulnerability - showcasing how semantic conflicts can generate humour without relying on sound patterns; and 3) the ironic juxtaposition of verbal affirmation against visible distress - illustrating how meaning-based conflicts drive the humour independently of linguistic devices like rhyme or wordplay.

This example illustrates how different humour mechanisms (self-reference, alliteration, and semantic conflict) can function independently, each contributing uniquely to the joke’s overall impact.

Pun usage demonstrated selective associations, showing moderate correlations with sound-based features including rhyme (Pearson: 0.29, Spearman: 0.26) and homonyms (Pearson: 0.28, Spearman: 0.20). However, correlations with other mechanisms were weaker (Pearson(r) <0.25, Spearman (r) <0.10). This pattern suggests that while puns integrate sound patterns and multiple meanings, they operate through relatively distinct linguistic pathways in humour generation.

### **Mechanism Correlation Conclusion**

The comparison between Pearson and Spearman correlations reveals that humour mechanisms interact through multiple parallel pathways rather than a single unified system. While some mechanisms show strong linear relationships in extreme cases (as evidenced by high Pearson correlations), the lower Spearman correlations suggest these relationships are not typical of most humorous content. We observe the following:

- Outlier-driven relationships: Particularly between semantic conflicts and rhyme
- Stable moderate correlations: Especially in homonym-based features
- Independent mechanisms: Such as alliteration and self-reference
- Context-dependent interactions: Varying strength of relationships based on humour style and content

This analysis suggests that effective humour can be constructed through various combinations of these mechanisms, with no single pattern dominating across all instances.

## **4.3 Affective and Emotional Patterns**

Examination of affective patterns revealed distinct emotional signatures and sentiment characteristics across humour styles, providing insights into their underlying psychological mechanisms.

### *4.3.1 Distribution of Primary Emotions*

Table 2 presents the frequency distribution of primary emotions across different humour styles. Each style demonstrated characteristic patterns in emotional expression.

Emotion Distribution						
	Anger	Fear	Joy	Love	Sadness	Surprise
<b>Affiliative</b>	18	8	16	3	3	1
<b>Aggressive</b>	28	7	13	1	8	0
<b>Neutral</b>	34	4	39	1	0	2
<b>Self-deprecating</b>	14	4	15	0	11	2
<b>Self-enhancing</b>	15	3	30	4	7	2

Table 2: Emotion Distribution by Humour Style

**Self-Enhancing Humour.** Self-enhancing humour demonstrated clear joy ( $n = 30$ ) predominance with positive sentiment polarity (0.227) and the highest confidence in classification (0.889) among all styles. For instance: *“Self-love is the greatest gift you can give yourself, besides a lifetime supply of cheese”*. This example shows maximum positive sentiment polarity (1.0) and joy (0.978), typical of self-enhancing humour’s optimistic orientation. The secondary emotion associated with self-enhancing is anger ( $n=15$ ) an example of a self-enhancing joke associated with anger is: *“The question isn’t who is going to let me, it’s who is going to stop me”*. This example has anger (0.95) and a positive sentiment (0.99).

**Aggressive Humour.** This showed the highest frequency of anger ( $n = 28$ ) among all styles, with negative sentiment polarity (-0.039). The presence of joy-related content ( $n = 13$ ) suggests complex emotional dynamics within this style. Examples of aggressive humour with anger and joy emotions:

- **Anger (0.998):** *“A man wakes from a coma. His wife changes out of her black clothes and, irritated, remarks, “I really cannot depend on you in anything, can I!”*
- **Joy (0.997):** *“ I’d like to leave you with one thought, but I’m not sure you have a place to put it!”*

**Self-Deprecating Humour.** Analysis revealed that self-deprecating humour maintained the most balanced emotional distribution, with no significant differences between positive and negative emotion frequencies. The distribution showed comparable levels of joy ( $n = 15$ , 32.6%), anger ( $n = 14$ , 30.4%), and sadness ( $n = 11$ , 23.9%), with near-neutral sentiment polarity (mean = -0.002, SD = 0.098). Examples of self-deprecating humour with anger, joy and sadness emotions:

- **Anger (0.991):** *“You are offered \$50000 but if you accept it, the person you hate most in the entire world gets \$100000. Are you taking it? Yes. Why wouldn’t I want \$150000”*
- **Joy (0.997):** *“My memory is as reliable as a goldfish’s. I can forget my own name while introducing myself to someone”*
- **Sadness (0.947):** *“I can’t deny that I made a lot of mistakes when I was younger. I’m older now, so I can make different yet more severe mistakes”*

**Affiliative Humour.** Affiliative humour demonstrated a relatively even distribution between anger (18) and joy (16), with positive overall polarity (0.145), suggesting its role in social bonding through shared emotional experiences. Examples of affiliative humour with anger and joy emotions with positive sentiments:

- **Anger ( 0.837):** *“Two cows are in a field. The first one says, “Mooooo!” The second one replies, “that’s what I was going to say!”*
- **Joy (0.549):** *“Money won’t buy happiness, but it will pay the salaries of a large research staff to study the problem”*



**Neutral Content.** Neutral content exhibited high frequencies of both joy ( $n = 39$ , 48.8%) and anger ( $n = 34$ , 42.5%), while maintaining significantly lower subjectivity scores (mean = 0.273, SD = 0.089) compared to all humour styles. An example of neutral category with joy emotion, positive sentiment and subjectivity of 0.3: *“Buying guide: find the best outdoor patio umbrella for your home”*

#### 4.3.2 Sentiment and Confidence Patterns

Table 3 presents the distribution of confidence scores and affective metrics across humour styles. Analysis revealed systematic variations in sentiment characteristics and classification confidence across different styles.

Style	N	Confidence	Polarity	Subjectivity	Sarcasm (%)
Affiliative	49	0.748	0.145	0.406	18.4
Self-enhancing	61	0.889	0.227	0.430	32.8
Self-deprecating	46	0.811	-0.002	0.474	8.7
Aggressive	57	0.780	-0.039	0.404	29.8
Neutral	80	0.887	0.089	0.273	48.8

Table 3: Distribution of Confidence Scores and Affective Metrics Across Humour Styles

**Sentiment Polarity.** Analysis revealed a clear demarcation in sentiment polarity across styles. Self-enhancing and affiliative humour maintained positive polarity (0.227 and 0.145 respectively), while aggressive humour exhibited negative polarity (-0.039). Self-deprecating humour demonstrated near-neutral polarity (-0.002), suggesting balanced emotional content. Neutral content showed mild positive polarity (0.089), positioned between the extremes of other styles.

**Subjectivity Patterns.** Subjectivity analysis revealed distinct patterns across styles. Self-deprecating humour exhibited the highest subjectivity (0.474), indicating strong emotional investment in content. Conversely, neutral content showed markedly lower subjectivity (0.273), consistent with its more objective nature. Affiliative and aggressive styles showed comparable subjectivity levels (0.406 and 0.404 respectively).

**Classification Confidence.** Model confidence demonstrated systematic variation across styles. Self-enhancing humour and neutral content showed the highest confidence scores (0.889 and 0.887 respectively), suggesting more distinctive linguistic and emotional patterns. Affiliative humour exhibited the lowest confidence (0.748), indicating greater classification uncertainty for this style.

**Sarcasm Distribution.** Sarcasm presence varied substantially across styles, with neutral content showing the highest frequency (48.8% of instances) and self-deprecating humour the lowest (8.7%). This distribution suggests that sarcasm may serve different functions across humour styles, potentially contributing to style differentiation. Example of sarcastic jokes in the dataset:

- **Self-enhancing:** *“You are never too old to play. You are only too old for low-rise jeans.”*
- **Affiliative:** *“Makeup can only make you look pretty on the outside but it doesn’t help if you’re ugly on the inside. Unless you eat the makeup.”*
- **Self-deprecating:** *“I’m so indecisive, I couldn’t choose my favorite color even if you held a rainbow in front of me!”*

These patterns indicate that humour styles employ distinct emotional mechanisms and linguistic strategies. Self-enhancing and affiliative styles maintain predominantly positive emotional valence, while aggressive and self-deprecating styles demonstrate more complex emotional patterns. The variation in subjectivity and sarcasm levels further differentiates these styles, with self-deprecating humour showing heightened emotional investment and neutral content maintaining greater objectivity. While these patterns reveal distinct characteristics of different humour styles, the classification process faces several challenges, which we examine in detail in the following section on error analysis.

#### 4.4 Error Analysis and Misclassification Patterns

Systematic analysis of classification errors revealed distinct patterns in model behaviour, with a predominant challenge in affiliative humour classification. This finding aligns with previously identified performance limitations in the single-model approach [Kenneth Ogbuka et al., 2024].

##### 4.4.1 Primary Misclassification Categories

Analysis of error patterns revealed three dominant misclassification types, each characterised by distinct feature combinations:

**Affiliative–Neutral Confusion (n = 8).** These misclassifications demonstrated moderate classification confidence (mean = 0.649) with distinctive sentiment characteristics: elevated polarity (0.197) combined with reduced subjectivity (0.296). In particular, these cases exhibited minimal semantic conflicts (mean = 9.500), suggesting potential oversimplification of humorous content. An example affiliative joke misclassified as neutral is: *Joke: “There are only 3 things that tell the truth: 1 - Young Children 2 - Drunks 3 - Leggings”*.

**Affiliative–Self-enhancing Confusion (n = 8).** Cases in this category showed moderate confidence levels (mean = 0.709) with mild positive sentiment polarity (0.094). A distinguishing feature was the elevated presence of exaggeration (mean = 1.375), indicating potential confusion between social and self-directed humour mechanisms.

**Affiliative–Aggressive Confusion (n = 7).** These cases exhibited the lowest confidence scores (M = 0.611) and negative sentiment polarity (-0.037). The high semantic conflict count (M = 29.857) suggests that complex linguistic structures may contribute to classification ambiguity.

Detailed examples of misclassified or confused categories are presented in subsection 4.4.3.

##### 4.4.2 Feature-Specific Error Analysis

Analysis of classification errors revealed systematic patterns across multiple dimensions, as detailed in Table 4.

**Confidence Patterns.** Model confidence demonstrated significant variation across error types ( $p < 0.0001$ ), with distinct patterns emerging:

- Affiliative–aggressive misclassifications exhibited the lowest confidence (mean = 0.611, SD = 0.179).
- Self-deprecating misclassifications showed unexpectedly high confidence (mean = 0.811, SD = 0.160).

**Sentiment Characteristics.** Sentiment analysis revealed significant associations between misclassification types and polarity patterns ( $p = 0.0001$ ):

- Affiliative humour misclassified as aggressive showed negative polarity (-0.037), deviating from typical affiliative patterns.

- Self-deprecating humour misclassified as affiliative exhibited unexpected positive sentiment (0.403).
- Subjectivity scores varied systematically with error types, particularly in cases involving self-deprecating humour.

**Semantic Structure.** Semantic conflict analysis revealed significant patterns ( $p = 0.0003$ ):

- The presence of high semantic conflict counts ( $>25$ ) strongly influenced misclassifications.
- Affiliative–aggressive confusions showed elevated semantic conflicts (mean = 29.857, SD = 63.006).
- Aggressive–affiliative misclassifications demonstrated the highest semantic conflict counts (mean = 37.833, SD = 74.778).

Error Type	Confidence		Semantic Conflicts		Sentiment	
	Mean	SD	Mean	SD	Polarity	Subj.
<i>Affiliative Misclassifications</i>						
→ Aggressive	0.611	0.179	29.857	63.006	-0.037	0.304
→ Neutral	0.649	0.226	9.500	7.309	0.197	0.296
→ Self-deprecating	0.776	0.289	22.000	15.620	0.175	0.510
→ Self-enhancing	0.709	0.169	6.500	6.908	0.094	0.256
<i>Aggressive Misclassifications</i>						
→ Affiliative	0.663	0.178	37.833	74.778	-0.024	0.568
→ Neutral	0.746	0.171	7.600	8.792	0.016	0.150
→ Self-deprecating	0.628	0.227	10.500	9.192	0.300	0.442
→ Self-enhancing	0.698	0.268	0.500	0.707	0.356	0.750
<i>Self-deprecating Misclassifications</i>						
→ Affiliative	0.775	0.074	17.333	11.015	0.403	0.550
→ Aggressive	0.811	0.160	14.000	7.937	-0.046	0.554
→ Neutral	0.773	0.159	5.500	7.778	0.000	0.275
→ Self-enhancing	0.881	–	28.000	–	0.250	0.250
<i>Self-enhancing Misclassifications</i>						
→ Affiliative	0.587	0.423	7.000	4.243	-0.250	0.500
→ Aggressive	0.745	–	0.000	–	0.062	0.167
→ Neutral	0.480	0.026	11.500	2.121	0.000	0.000
→ Self-deprecating	0.688	0.264	6.800	7.791	-0.076	0.383

Table 4: Characteristics of Misclassification Types

**Additional Feature Patterns.** Analysis of supplementary features revealed:

- Confidence scores generally decreased when the model confused similar styles (e.g, affiliative and aggressive).
- Sentiment polarity often deviated significantly from expected patterns for specific humour styles.
- Higher semantic conflict counts corresponded with increased likelihood of misclassification between affiliative and aggressive styles.
- Sarcasm probability showed significant variation across error types, suggesting its role in misclassifications.

These patterns indicate that misclassifications arise from complex interactions between linguistic and affective features, particularly when these features deviate from style-typical patterns. The systematic nature of these deviations suggests specific weaknesses in the model’s ability

to distinguish between stylistically similar humour categories, especially in cases with high semantic complexity.

#### 4.4.3 Analysis of Representative Misclassifications

Detailed examination of misclassified instances revealed systematic patterns in classification errors. We selected eight representative cases (two from each humour style) to illustrate key error patterns, focusing on cases with the most distinctive feature combinations.

##### Affiliative Humour Misclassifications.

###### 1. Error Type: Affiliative misclassified as Self-deprecating:

- *Joke*: “As best man, it is my job to tell you about the groom, and all the embarrassing things that have happened to him in the 28 years leading up to what was the happiest day of his life until I started this speech”
- confidence score: 0.447
- Top features: “embarrassing”: 0.30, “my”: 0.17, “groom”: -0.17, “i”: 0.12, “28”: -0.10
- Linguistic features: High semantic conflict count (40), self-reference count (2 instances)
- Affective Profile: Sarcasm (false), negative sentiment (0.989 confidence), subjectivity (0.3), emotion (sadness—0.932 confidence)
- Target: self-targeted: true, other-targeted: true, situation-targeted: false
- Analysis: The model focusing on the word “embarrassing” and personal references “my”, and the strong negative sentiment and emotion, and high semantic conflict in the sentence triggered self-deprecating classification despite the social bonding context typical of wedding speeches.

###### 2. Error Type: Affiliative misclassified as Aggressive:

- *Joke*: “What did one DNA say to the other DNA? these genes make me look fat”
- confidence score: 0.781
- Top features: “me”: 0.05, “say”: -0.03, “genes”: -0.03, “one”: -0.02, “did”: -0.02
- Linguistic features: semantic conflicts (19), self-reference count (1)
- Affective Profile: Sarcasm probability (0.999), negative sentiment (0.99 confidence), subjectivity (0.375), emotion (joy—0.865 confidence)
- Target: self-targeted: true, other-targeted: false, situation-targeted: false
- Analysis: The model’s high confidence (0.781) in the aggressive classification appears driven by the sentence’s sarcastic nature, negative sentiment, and homophones “genes/jeans” that could be interpreted as mockery.

##### Aggressive Humour Misclassifications.

###### 1. Error Type: Aggressive misclassified as Affiliative:

- *Joke*: “After every sentence I say you say ketchup and rubber buns. what did you eat for breakfast? “ketchup & rubber buns.” what did you eat for lunch? “ketchup & rubber buns.” what did you eat for dinner? “ketchup & rubber buns.” what do you do when you see a hot girl? “ketchup & rubber buns.” YOU WHERE RUBBING MY GF’S WHAT?!?!”
- confidence score: 0.42
- Top features: “every”: 0.08, “I”: 0.06, “say”: -0.04, “sentence”: 0.03, “MY”: 0.03
- Linguistic features: semantic conflicts (190), self-reference count (2), rhyme count (14)

- Affective Profile: Sarcasm (false), positive sentiment (0.99 confidence), subjectivity (0.85), emotion (love - 0.98 confidence)
- Target: self-targeted: true, other-targeted: true, situation-targeted: false
- Analysis: The model's misclassification appears driven by the high positive sentiment, high positive emotion of love, and playful repetition (high rhyme count) rather than the hostile twist ending.

## 2. Error Type: Aggressive misclassified as Self-enhancing:

- Joke: *Cats have nine lives. Makes them ideal for experimentation*
- confidence score: 0.51
- Top features: "experimentation": 0.03, "nine": -0.02, "have": -0.01, "Cats": -0.009, "Makes": -0.009
- Linguistic features: semantic conflicts (1), self-reference count (2)
- Affective Profile: Sarcasm probability (1.0), positive sentiment (0.99 confidence), subjectivity (1.0), emotion (joy - 0.99 confidence)
- Target: self-targeted: false, other-targeted: false, situation-targeted: true
- Analysis: The model's misclassification appears driven by the high positive sentiment, high positive emotion of joy, presence of sarcasm, and lack of typical aggressive markers like profanity or direct attacks.

## Self-deprecating Humour Misclassifications.

### 1. Error Type: Self-deprecating misclassified as Self-enhancing:

- Joke: *"You're guessing that out of the 8 billion people here on Earth, I'm going to chase someone who doesn't even like me? Well, watch me closely, because that's exactly what I'm going to do"*
- confidence score: 0.88
- Top features: "me": 0.12, "Well": -0.09, "chase": -0.08, "I": 0.08, "You": -0.07
- Linguistic features: semantic conflicts (28), self-reference count (4)
- Affective Profile: Sarcasm (false), negative sentiment (0.99 confidence), subjectivity (0.25), emotion (joy - 0.65 confidence)
- Target: self-targeted: true, other-targeted: true, situation-targeted: false
- Analysis: The high self-reference makes the joke to be likely classified as either self-deprecating or self-enhancing. Even though the sentiment was negative, the misclassification arises from the emotion being joy, which conflicts with the sentiment and the high semantic conflicts of 28 reflect this.

### 2. Error Type: Self-deprecating misclassified as Aggressive:

- Joke: *"What would have happened if you exterminated the ugliest guy and the dumbest guy in the world yesterday? Right, this post wouldn't exist."*
- confidence score: 0.99
- Top features: "dumbest": -0.02, "guy": 0.01, "would": -0.008, "wouldn't": -0.005, "Right": 0.005
- Linguistic features: semantic conflicts (20), self-reference count (0)
- Affective Profile: Sarcasm probability (1.0), negative sentiment (0.99 confidence), subjectivity (0.55), emotion (anger - 0.65 confidence)
- Target: self-targeted: false, other-targeted: true, situation-targeted: false
- Analysis: Negative sentiment and emotions are both highly correlated with self-deprecating and aggressive humour; however, the lack of self-reference and the presence of others-reference led the model to mislabel.

## Self-enhancing Humour Misclassifications.



### 1. Error Type: Self-enhancing misclassified as Affiliative:

- Jokes: *“I was gonna tell a joke about pizza, but it’s too cheesy”*
- confidence score: 0.29
- Top features: “I”: 0.21, “pizza”: -0.18, “joke”: -0.14, “cheesy”: 0.12, “tell”: -0.09
- Linguistic features: semantic conflicts (10), self-reference count (1)
- Affective Profile: Sarcasm (false), negative sentiment (0.99 confidence), subjectivity (1.0), emotion (joy - 0.90 confidence)
- Target: self-targeted: true, other-targeted: false, situation-targeted: false
- Analysis: The low confidence score (0.29) reflects uncertainty, indicating model difficulty. The self-referential context (“I”) and playful tone suggest self-enhancing humour. However, the light-hearted nature, lack of personal boasting, and emphasis on social food humour, coupled with high joy emotion (0.90), likely caused misclassification as affiliative humour.

### 2. Error Type: Self-enhancing misclassified as Self-deprecating:

- Joke: *“I went to the doctor with a strawberry growing on my bum. The doctor said I’ve got some cream for that”*
- confidence score: 0.33
- Top features: “strawberry”: -0.20, “I”: 0.19, “growing”: -0.16, “bum”: 0.10, “said”: -0.10
- Linguistic features: semantic conflicts (4), self-reference count (3)
- Affective Profile: Sarcasm (false), negative sentiment (0.99 confidence), subjectivity (0.0), emotion (anger - 0.54 confidence)
- Target: self-targeted: true, other-targeted: false, situation-targeted: false
- Analysis: The misclassification probably resulted from the overlap of humour characteristics. The joke’s self-referential nature (3 instances) and mildly embarrassing scenario align with self-deprecating humour. However, the speaker’s light-hearted tone and ability to laugh at a personal situation suggest self-enhancing intent. Low semantic conflict (4) and the unexpected association of “anger” as the dominant emotion further confused the model.

## Error Pattern Analysis

Analysis of the misclassified instances revealed five fundamental challenges in humour style classification:

1. **Emotional Ambiguity:** The model demonstrates significant sensitivity to conflicting emotional signals. This is particularly evident in self-enhancing humour, where positive sentiment often co-occurs with contradictory emotional markers (e.g., anger or fear), leading to misalignment between detected and intended emotional content.
2. **Context Misinterpretation:** The model shows systematic difficulties in accurately interpreting social contexts, particularly in discriminating between self-directed and other-directed humorous content. This limitation suggests insufficient incorporation of broader contextual cues in the classification process.
3. **Self-Reference Confusion:** Self-referential content creates particular classification challenges, most notably in distinguishing between self-deprecating and self-enhancing styles. This suggests potential limitations in the model’s ability to interpret the valence of self-referential statements.
4. **Target Ambiguity:** The model exhibits reduced accuracy in cases with multiple potential targets (self, others, or situations). This limitation is particularly pronounced in instances

where humorous content contains overlapping targeting mechanisms.

5. **Feature Interference:** The simultaneous presence of multiple humour mechanisms (e.g., wordplay, sarcasm, semantic conflicts) creates competing classification signals, suggesting limitations in the model's ability to integrate multiple linguistic features.

These patterns demonstrate the effectiveness of the XAI framework in identifying both feature-level drivers of classification decisions and systematic classification challenges. The findings suggest that improving model performance will require enhanced handling of cases involving multiple interacting humour mechanisms, particularly for affiliative humour where the interplay between emotional and linguistic signals exhibits greater complexity.

## V CONCLUSION AND FUTURE WORK

This study introduces a comprehensive XAI framework for understanding humour style classification, revealing significant patterns in how linguistic mechanisms and emotional features interact across humour styles. Social forms of humour, particularly affiliative and aggressive styles, demonstrate higher rates of wordplay and semantic conflicts. The emotional signatures vary markedly between styles, with self-enhancing humour showing consistently positive polarity (0.227) and high confidence (0.889), while aggressive humour tends toward negative polarity (-0.039). Our correlation analyses reveal complex relationships between linguistic mechanisms, particularly in the relationship between semantic conflicts and rhyme, where Pearson correlation ( $r = 0.95$ ) suggests a strong linear relationship, but the much lower Spearman correlation ( $r = 0.32$ ) indicates this relationship is heavily influenced by outliers rather than being typical across most instances.

The classification challenges highlight important considerations for future development. The model shows varying confidence levels across styles, particularly struggling with affiliative humour (lowest average confidence: 0.748). These challenges often stem from emotional ambiguity and context misinterpretation, especially when multiple humour mechanisms interact. Self-reference and target identification emerged as critical factors in classification decisions, suggesting their importance for future model development.

The practical implications of our findings extend beyond theoretical understanding. Our framework provides actionable insights for researchers studying humour's role in psychological well-being and social communication, while offering a valuable template for applying XAI to other subjective classification tasks in digital humanities. The framework's ability to explain model decisions enhances transparency and trust in automated humour analysis, supporting applications in mental health assessment, content moderation, and digital humanities research.

Despite the valuable insights provided by this study, several important limitations must be acknowledged. Our analysis focuses on a single model configuration (ALI+XGBoost) with a dataset of 1,463 English-language instances, which inherently limits the generalisability of our findings. The feature detection components faced significant technical constraints, particularly in terms of dictionary coverage and algorithmic accuracy. The CMU Pronouncing Dictionary covered 90.75% of unique words in our dataset for rhyme detection, while WordNet synset coverage reached 88.50% of the vocabulary, leaving significant gaps in handling slang and neologisms. The alliteration detector showed particular limitations, such as considering only the first phoneme, leading to false positives like grouping words beginning with different sounds (e.g., "when/one"), counting repeated words as alliteration, and grouping words regardless of their proximity in the text. This contributed to moderate correlations with homonym detection (Pearson: 0.21, Spearman: 0.27), which is partially due to overlapping feature detection rather

than true linguistic relationships.

A significant methodological constraint lies in our reliance on pre-trained models for sarcasm (60.7% accuracy), sentiment (93.2% accuracy), and emotion detection (93.8% accuracy), not specifically trained on humorous content. Humorous text contains unique linguistic patterns, irony, and complex emotional layers that may not be well-captured by these general-purpose models. The binary classification of features, such as sarcasm detection, oversimplifies complex linguistic phenomena that exist on a spectrum. Additionally, automated sentiment and emotion detection may not fully capture the subtle, often contradictory emotional cues present in different humour styles, like simultaneous positive and negative emotions in self-deprecating humour.

Several analytical constraints affect our findings. The study does not address temporal or contextual variations in humour interpretation, and the analysis of misclassifications may not capture all possible error patterns due to the limited sample size. Future research should expand to multiple models and larger, more diverse datasets, incorporating cross-cultural perspectives and multilingual analysis. Enhanced feature detection algorithms are needed, particularly for complex linguistic features like multi-word puns and contextual wordplay, along with better handling of informal language. The development of specialised models trained specifically on humorous content would improve the accuracy of sentiment and emotion detection in comedic contexts. Additionally, more refined methods for capturing contextual and cultural factors would enhance the framework's applicability across different contexts, while integration of temporal analysis would provide deeper insights into how humour interpretation varies over time and across different social settings. These improvements would strengthen the framework's reliability and broaden its applicability in real-world scenarios.

## VI ACKNOWLEDGEMENT

This research was supported by the Petroleum Technology Development Fund (PTDF) of Nigeria.

## References

- Muhammad Abulaish and Ashraf Kamal. Self-Deprecating Sarcasm Detection: An Amalgamation of Rule-Based and Machine Learning Approach. In *Proceedings - 2018 IEEE/WIC/ACM International Conference on Web Intelligence, WI 2018*, pages 574–579. Institute of Electrical and Electronics Engineers Inc., 1 2018. ISBN 9781538673256. doi: 10.1109/WI.2018.00-35.
- Usman Ahmed, Rutvij H. Jhaveri, Gautam Srivastava, and Jerry Chun-Wei Lin. Explainable Deep Attention Active Learning for Sentimental Analytics of Mental Disorder. *ACM Transactions on Asian and Low-Resource Language Information Processing*, 7 2022. ISSN 2375-4699. doi: 10.1145/3551890.
- Arooba Amjad and Rabia Dasti. Humor styles, emotion regulation and subjective well-being in young adults. *Current Psychology*, 41(9):6326–6335, 9 2022. ISSN 19364733. doi: 10.1007/s12144-020-01127-y.
- Whitney Anderson and Nancy Di Tunnariello. Aggressive humor as a negative relational maintenance behavior during times of conflict. *Qualitative Report*, 21(8):1513–1530, 8 2016. ISSN 21603715. doi: 10.46743/2160-3715/2016.2149.
- Issa Annamoradnejad and Gohar Zoghi. CoBERT: Using BERT Sentence Embedding in Parallel Neural Networks for Computational Humor. 4 2020. URL <http://arxiv.org/abs/2004.12765>.
- Guo Hai Chen and Rod A. Martin. A comparison of humor styles, coping humor, and mental health between Chinese and Canadian university students. *Humor*, 20(3):215–234, 8 2007. ISSN 09331719. doi: 10.1515/HUMOR.2007.011.
- Yuyan Chen, Yichen Yuan, Panjun Liu, Dayiheng Liu, Qinghao Guan, Mengfei Guo, Haiming Peng, Bang Liu, Zhixu Li, and Yanghua Xiao. Talk Funny! A Large-Scale Humor Response Dataset with Chain-of-Humor Interpretation. In *The Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI-24)*, pages 17826–17834, 2024. URL <https://ojs.aaai.org/index.php/AAAI/article/view/29736>.

- Kounteyo Roy Chowdhury, Arpan Sil, and Sharvari Rahul Shukla. Explaining a Black-Box Sentiment Analysis Model with Local Interpretable Model Diagnostics Explanation (LIME). In *Communications in Computer and Information Science*, volume 1440 CCIS, pages 90–101. Springer Science and Business Media Deutschland GmbH, 2021. ISBN 9783030814618. doi: 10.1007/978-3-030-81462-5{\\_}9.
- Karina Cortinas-Lorenzo and Gerard Lacey. Toward Explainable Affective Computing: A Review. *IEEE Transactions on Neural Networks and Learning Systems*, 2024. ISSN 21622388. doi: 10.1109/TNNLS.2023.3270027.
- Victor De Marez, Thomas Winters, and Ayla Rigouts Terryn. THInC: A Theory-Driven Framework for Computational Humor Detection. In *International Workshop on Artificial Intelligence and Creativity*, Santiago de Compostela,, 9 2024. URL <http://arxiv.org/abs/2409.01232>.
- Abbas Edalat. Self-initiated humour protocols: An algorithmic approach for learning to laugh. *PsyArXiv*, 5:1–14, 11 2023. doi: <https://doi.org/10.31234/osf.io/w9cvx>.
- Abbas Edalat, Ruoyu Hu, Zeena Patel, Neophytos Polydorou, Frank Ryan, and Dasha Nicholls. Self-Initiated Humour Protocol: A pilot study with an AI agent. *PsyArXiv*, 2024. doi: 10.31234/osf.io/f32xq.
- Jochen Hartmann, Mark Heitmann, Christian Siebert, and Christina Schamp. More than a Feeling: Accuracy and Application of Sentiment Analysis. *International Journal of Research in Marketing*, 40(1):75–87, 3 2023. ISSN 01678116. doi: 10.1016/j.ijresmar.2022.05.005.
- Ashraf Kamal and Muhammad Abulaish. Self-deprecating Humor Detection: A Machine Learning Approach. In Le-Minh Nguyen, Xuan-Hieu Phan, Kôiti Hasida, and Satoshi Tojo, editors, *Computer Linguistics*, volume 1215 of *Communications in Computer and Information Science*, pages 483–484. Springer Singapore, Singapore, 2020. ISBN 978-981-15-6167-2. doi: 10.1007/978-981-15-6168-9. URL <http://link.springer.com/10.1007/978-981-15-6168-9>.
- Przemysław Kazienko, Julita Bielaniec, Marcin Gruza, Kamil Kanclerz, Konrad Karanowski, Piotr Miłkowski, and Jan Kocoń. Human-centered neural reasoning for subjective content processing: Hate speech, emotions, and humor. *Information Fusion*, 94:43–65, 6 2023. ISSN 15662535. doi: 10.1016/j.inffus.2023.01.010.
- Mary Ogbuka Kenneth, Foad Khosmood, and Abbas Edalat. Systematic Literature Review: Computational Approaches for Humour Style Classification. Technical report, 2024. URL <https://www.semanticscholar>.
- Mary Kenneth Ogbuka, Foad Khosmood, and Abbas Edalat. A Two-Model Approach for Humour Style Recognition. In *Proceedings of the 4th International Conference on Natural Language Processing for Digital Humanities*, pages 259–274, Miami, 11 2024. Association for Computational Linguistics. URL <https://aclanthology.org/2024.nlp4dh-1.25/>.
- Nicholas Kuiper, Gillian Kirsh, and Nadia Maiolino. Identity and Intimacy Development, Humor Styles, and Psychological Well-Being. *Identity*, 16(2):115–125, 4 2016. ISSN 1532706X. doi: 10.1080/15283488.2016.1159964.
- Scott Lundberg and Su-In Lee. A Unified Approach to Interpreting Model Predictions. In *31st Conference on Neural Information Processing Systems*, Long Beach, 5 2017. URL <http://arxiv.org/abs/1705.07874>.
- Qing Lyu, Marianna Apidianaki, and Chris Callison-Burch. Towards Faithful Model Explanation in NLP: A Survey. *Computational Linguistics*, pages 657–723, 2024. doi: 10.1162/coli. URL <https://doi.org/10.1162/coli>.
- Rutal Mahajan and Mukesh Zaveri. An automatic humor identification model with novel features from Berger’s typology and ensemble models. *Decision Analytics Journal*, 11, 6 2024. ISSN 27726622. doi: 10.1016/j.dajour.2024.100450.
- Rowan Mann and Tomislav Mikulandric. CLEF 2024 JOKER Tasks 1-3: Humour identification and classification. In *Conference and Labs of the Evaluation Forum*, Grenoble, 9 2024.
- Rod A Martin and Thomas Ford. *The psychology of humor: An integrative approach*. Academic Press, 2nd edition, 2018.
- Rod A Martin, Patricia Puhlik-Doris, Gwen Larsen, Jeanette Gray, and Kelly Weir. Individual differences in uses of humor and their relation to psychological well-being: Development of the Humor Styles Questionnaire. *Journal of Research in Personality*, 37:48–75, 2003. URL [www.elsevier.com/locate/jrp](http://www.elsevier.com/locate/jrp).
- Reynier Ortega-Bueno, Paolo Rosso, and José E. Medina Pagola. Multi-view informed attention-based model for Irony and Satire detection in Spanish variants. *Knowledge-Based Systems*, 235, 1 2022. ISSN 09507051. doi: 10.1016/j.knosys.2021.107597.
- Gabriel Ichcanziho Pérez-Landa, Octavio Loyola-González, and Miguel Angel Medina-Pérez. An explainable artificial intelligence model for detecting xenophobic tweets. *Applied Sciences (Switzerland)*, 11(22), 11 2021. ISSN 20763417. doi: 10.3390/app112210801.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. “Why Should I Trust You?” Explaining the Predictions of Any Classifier. In *Proceedings of NAACL-HLT*, pages 97–101, San Diego, 6 2016. Association for

Computational Linguistics. URL <https://github>.

Serkan Volkan Sari. Was it just joke? Cyberbullying perpetrations and their styles of humor. *Computers in Human Behavior*, 54:555–559, 1 2016. ISSN 07475632. doi: 10.1016/j.chb.2015.08.053.

Anindya Sen. Humour Analysis and Qualitative Research. Technical report, University of Surrey, 2012.

Donghai Zhang, Wei Song, Lizhen Liu, Chao Du, and Xinlei Zhao. Investigations in automatic humor recognition. In *Proceedings - 2017 10th International Symposium on Computational Intelligence and Design, ISCID 2017*, volume 1 2018-January, pages 272–275. Institute of Electrical and Electronics Engineers Inc., 7 2017. ISBN 9781538636749. doi: 10.1109/ISCID.2017.160.

Hong Zhu, Yilin Ou, and Zimeng Zhu. Aggressive humor style and cyberbullying perpetration: Normative tolerance and moral disengagement perspective. *Frontiers in Psychology*, 13, 12 2022. ISSN 16641078. doi: 10.3389/fpsyg.2022.1095318.

## A SUMMARY OF RELATED WORKS

Table 5 summarises the key works discussed in section 2 (Related works), highlighting the authors’ specific tasks, datasets used, features extracted, classification approaches, XAI methods employed, and their achieved results. This overview demonstrates both the progress made in computational humour analysis and explainability techniques while highlighting the gaps our work addresses.

Author	Task	Dataset	Extracted Features	Classifiers	XAI Method	Results
Pérez-Landa et al. [2021]	Xenophobia	Experts Xenophobia Database (EXD) (10,057 tweets) Pitropakis Xenophobia Database (PXD)(5814 tweets)	TF-IDF, Bag of words (BOW), Word 2 vector (W2V), Keywords, Sentiment, Emotion, and Syntactic	C45, KNN, Rusboost, UnderBagging.	Feature engineering (keywords, sentiment, emotions and syntactic)	F1: 76.8% (EXD) AUC: 0.864 (EXD) F1: 73.4% (PXD) AUC: 0.794 (PXD)
Chowdhury et al. [2021]	Sentiment	Twitter API (2-years period)	W2V, Sentiment Specific Word Embedding (SSWE), GloVe, FastText embeddings	Bi-directional LSTM	LIME	ACC: 72% F1: 72%
Ahmed et al. [2022]	Sentiment	Patient-authored text (15,044)	Fuzzy logic rules, GloVe, Cosine Similarity, Keywords	Bi-directional LSTM Attention with fuzzy classification	Fuzzy logic rules	F1: 89%
Ortega-Bueno et al. [2022]	Irony, Satire	IroSva'19 shared task (3000) Barbieri'15-es (10,000) Salas'17-mx (10,888) HAHA'19 shared task (6000)	BERT embedding, Multilingual Universal Sentence encoding, semantics, affective, incongruity, stylistic and structural	Multiview Attention LSTM	Feature engineering (semantics, affective, incongruity, stylistic and structural)	F1: 70.4% (IroSva'19) F1: 95.7% (Barbieri'15) F1: 96% (Salas'17) F1: 80.6% (HAHA'19)
Zhang et al. [2017]	Humour	16000 One-Liner (32,002)	Contextual knowledge, Subjectivity, affective polarity, ambiguity, incongruity	-	Feature engineering (Contextual knowledge, Subjectivity, polarity, ambiguity, incongruity)	F1: 85%
Mann and Mikulandric [2024]	Humour, Pun	CLEF 2024 JOKER Tasks	TF-IDF	Logistic regression, Naive Bayes (NB), SVM, MarianMTModel		F1: 83% (Pun) F1: 61% (Humour)
De Marez et al. [2024]	Humour	SemEval 2021 Task 7	Polarity, emotions, offense, subjectivity, hate, stance, ambiguity, adult language	GA2M	Feature engineering (Polarity, emotions, offense, subjectivity, hate, stance, ambiguity, adult language)	F1: 85%
Mahajan and Zaveri [2024]	Humour	Yelp reviews	Emotive, incongruity, intensity	Stacking-based ensemble, NB, SVM, MLP, majority-vote ensemble	Feature engineering (Berger's typology-based incongruity features)	F1: 72.57%
Chen et al. [2024]	Humour	Chain-of-humour dataset	word embeddings	Pre-trained language models (PLMs)	Chain-of-thought	F1: 87%
Abulaish and Kamal [2018]	Sarcasm	Twitter API (7 weeks tweet)	Self-referencing, hyperbolic, part-of-speech	Decision tree, Naive Bayes, Bagging		F1: 94%
Kamal and Abulaish [2020]	Humour	Pun of the Day (4,826), 1600 One-Liner (32,002), Twitter API (20,000)	Self-deprecating pattern, exaggeration, word embedding, ambiguity, interpersonal effect, phonetic style.	Random forest		F1: 62% - 87% across datasets
Kenneth Ogbuka et al. [2024]	Humour Styles	Humour styles dataset	Sentence embedding	NB, Randon forest, XGBoost, DistilBERT		F1: 78.6%

Table 5: Summary of Related Works