

Thesis

by Aaye Vesa

Submission date: 20-Jun-2026 02:25PM (UTC+0530)

Submission ID: 2986457215

File name: Arya_Thesis.pdf (2.63M)

Word count: 11865

Character count: 70862

M. Tech.
Thesis

A HYBRID DEEP LEARNING AND TRANSFORMER BASED MODEL FOR MULTIMODAL SARCASAM DETECTION

A HYBRID DEEP LEARNING AND TRANSFORMER BASED
MODEL FOR MULTIMODAL SARCASAM DETECTION

Patel Arya Jigneshkumar

May
2026

18
Thesis Submitted
in Partial Fulfillment of the Requirements for the
Degree of

MASTER OF TECHNOLOGY

in
Artificial Intelligence

by
Patel Arya Jigneshkumar
(2K24/AFI/08)

4
Under the supervision of
Dr. Sanjay Kumar
(Dept of Computer Science & Engineering)



Department of Computer Science And Engineering
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi 110042,
India
MAY, 2026

ABSTRACT

In the past ten years social media has transformed how we communicate online and also has seen an increase in the use of casual, vague, at times sarcastic writing. While sarcasm is the most complex one can interpret, and with words going against images is more complicated today's sentiment analysis systems. For the above purpose which is to improve the above we formulated a Multimodal Sarcasm Detection system in the present piece of work which is designed to detect sarcasm from the whole text and also visual elements.

The methods used were: RoBERTa (Robustly Optimized BERT Pretraining Approach) for embedding text semantically, a Bidirectional Long Term Short Term Memory (BiLSTM) network for processing sequential context data for words in sentences and a pre learned ResNet50 convolutional neural network (CNN) for features extraction at higher order of visual information. Here we present what we term Emotion Heads: special projection networks, designed for each modality, that map text and visual representation to a what we may call emotion aware feature space before fusing modalities together. In fact, this kind of architecture generates the same situation as it typically generates the model of what's characteristically sarcastic: the occurrence of a cross modal emotional incongruity.

The Multimodal Sarcasm Detection (MMSD) dataset contains social media image-text pairs, and experiments were run on this data set using a preprocessing pipeline that includes use of the subword tokenization process from the RoBERTa tokenizer; standardized normalization for images based on ImageNet; and trained for 3 epochs using the AdamW optimizer with a learning rate of $2e-5$. The best validation accuracy (96.36%) and the Best validation F1-score (96.37%) was determined through this training. The test accuracy and weighted F1-score in the test set were 93.97% and 94.00% respectively, respectively, for its tested performance. The single-modal baselines and the two previous-state-of-the-art multimodal approaches were considerably outperformed by the model. Additionally, the test-set ROC-AUC was greater than 0.95, indicating good discriminating power. The F1-scores for both classes are around 0.94 per class, showing that the model performs well without being biased for any particular class, a usual problem with class imbalance for sarcasm detection.

Contents

Acknowledgement.....	i
Candidate's Declaration.....	ii
Certificate.....	iii
Abstract.....	iv
Contents.....	v
List of Tables.....	vii
List of Figures.....	viii
Chapter 1 Introduction	1
1.1 Overview.....	1
1.2 Motivation.....	3
1.3 Problem Statement	5
1.4 Contributions	5
Chapter 2 Literature Review.....	7
2.1 Sarcasm Detection in Traditional Text-Based.....	7
2.2 Multimodal Sarcasm Detection.....	8
2.3 Deep Learning Architectures for Multimodal Processing	10
2.4 Current Challenges and Research Gap.....	10
2.5 Summary and Research Motivation.....	11
Chapter 3 Methodology.....	12
3.1 System Overview.....	12
3.2 Dataset Description.....	12
3.3 Data Preprocessing Pipeline.....	13
3.4 Model Architecture.....	15
3.5 Training Framework.....	18
3.6 Summary of Methodology.....	19
Chapter 4 Results and Discussion.....	20
4.1 Evaluation Metrics	20
4.2 Training and Validation Results.....	20
4.3 Test Set Performance.....	22
4.4 Comparative Analysis.....	24

4.5 Ablation Study	25
Chapter 5 Conclusion and Future Scope	27
5.1 Conclusion	27
5.2 Future Scope.....	28
5.3 Challenges and Limitations.....	29
5.4 Social and Practical Impact.....	30
Bibliography.....	31

8
List of Tables

Table 1.1 Text Emotions Distribution (Sarcastic vs Non-sarcastic).....	3
Table 1.2 Image Emotions Distribution (Sarcastic vs Non-sarcastic).....	4
Table 3.1 Dataset Description - MMSD Splits	13
Table 4.1 Training Performance per Epoch.....	20
Table 4.2 Best Validation Confusion Matrix (Epoch 1).....	21
Table 4.3 Test Set Confusion Matrix	22
Table 4.4 Overall Performance Summary Across Training Phases.....	24
Table 4.5 Comparison between State of the Art Methods on MMSD Dataset.....	24
Table 4.6 According to Study - Impact of Individual Architectural Components.....	24

List of Figures

Fig 1.1 Distribution of Emotion in Text and Images.....	4
Fig 3.1 Proposed Model Architecture.....	15
Fig 4.1 Per-Class F1 Scores for the model	23
Fig 4.2 ROC Curve of the model.....	23

Chapter 1

INTRODUCTION

1.1 Overview

Sarcasm is one of the most complex and situationally specific types of Deliberate expression of meaning through its: human communication opposite. This linguistic phenomenon is heavily toned and context-dependent and shared. The interaction and understanding shared by communicators have become more common in digital. communication environments. With the rapid growth of social media platforms, messaging applications, and online forums, detecting and understanding sarcastic expressions in Cognitive science and technology today. One of the most pressing challenges in affective cognitive science and technology today has arisen from computational systems. Computing and natural language processing.

With the advent of human communication onto the digital landscape, people's understanding, and approach to communicating sarcasm has transformed. When face-to-face, before digital, there were plenty of contextual clues, like facial expressions, tone of voice and body language, which would easily convey sarcasm to human listeners. However, in an online environment these time-tested "hints" are lacking, and the user must resort to other means for sending sarcastic messages. Users have begun to employ two or more modalities in digital communication, knowing that by using a combination of textual messages with visual messages such as images, memes, GIF and emojis, they are adding an intensive layer of meaning that can effectively convey the intended message, even when delivered in a sarcastic manner.

This evolution has created new challenges for computational systems to understand communications intentions that involve human interaction. Humans can easily detect the sarcasm in a conversation as they grasp the meaning from context and emotional understanding, but AI systems have had difficulty with the subtle nuances and understated meanings found within sarcastic statements. This poses a particular difficulty in multimodal contexts where the sarcastic intent may result from semantic and affective interactions between lengthy textual content and images, where both text and image have to be understood individually, but also as a combination of text and image and image and text.

In this thesis, existing multimodal emotion recognition systems are analysed and a hybrid deep learning emotion recognition system, named the Multimodal Emotion BiLSTM Architecture, is presented to solve the limitations found in these emotion recognition systems. It combines the capabilities of cutting-edge neural architectures such as the RoBERTa architecture for in-depth semantic analysis of textual content, BiLSTM for sequence and contextual understanding, and ResNet50 to analyse the wide range of message emotions from images. These models are integrated with the architecture achieving improved accuracy and robustness of emotion classification when combined.

A unique feature of the proposed framework is introduction of dedicated Emotion Heads in various modalities. The multimodal fusion and final classification relies on a representation space learned from the Emotion Heads, which convert textual and visual information to the emotional space in advance. This method makes the system focus on the most emotional information rather than ignore other pieces of information, resulting in better feature representation discrimination and better overall performance in the multimodal emotion recognition system.

1.1.1 Research Novelty and Breakthrough Contribution

This thesis tackles the fundamental problem of multi-modal sarcasm detection, and introduces a Hybrid RoBERTa-BiLSTM and ResNet Architecture for the detection task. In previous works, more complex feature concatenation or graph networks were used, or more complex emotion integration was applied; in this work, a simple yet very effective framework is explicitly presented where the detection of emotion incongruity is regarded as one of the core detection mechanisms.

The following are the key technical innovations that differentiate this work from previous :

1. Transformer-Based Semantic Encoding (RoBERTa) : The first design decision made was to switch from BERT to RoBERTa. The distinction between the two is the way in which they learn when they are pre-trained: BERT uses the same masking pattern on each sentence that it trains on, whereas RoBERTa randomly shuffles the masking patterns across epochs, making the model learn true "word representations" instead of memorizing fixed patterns. The model obtains additional depth of understanding of sarcasm laced phrasing with training on a much larger corpus. Instead of trying to tweak all 12 layers on the task (which would be computationally expensive and have the potential to undo useful knowledge from the general purpose layers), the first eight layers remain frozen and only the last four layers are updated. This maintains the bulk of the load on the pretrained weights with the added layers capturing the individual linguistic nuances of sarcastic text.
2. Sequential Context Modeling (BiLSTM) : Transformers are great at seeing the connections between words in the whole sentence, but sarcasm works differently it sets the stage and then delivers a "punch" or contradiction in the middle. This is overlooked in a single pass of global attention since it appears to deal with positions somewhat symmetrically. A BiLSTM over the entire sentence of RoBERTa token embeddings directly tackles this by passing the sentence forward and the sentence in reverse through the network.
3. Deep Visual Feature Extraction (ResNet50) : Much of the sarcasm in multimodal texts is conveyed solely through the image: the text may be sincere by itself, but the presence of the image indicates irony. ResNet50 addresses this aspect of the issue. It has residual connections that enable deep training without gradient degradation and features are abstracted into concepts (e.g., visual sentiment, compositional contrast, mood of the scene) by the final pooling layer. These are the exact signifiers that tell you when an image doesn't go with the caption; the kind of mismatch that characterizes visually-driven sarcasm.

4. **Emotion-Specific Projection Heads:** Here, architecture diverges most significantly from the previous work. Instead of the shared fusion layer inferring the emotion from the combination of text and image, a separate projection head is used for the text representation and the image representation, before they are fused. The heads are small, feed-forward networks that compress the information contained in their input into 128 dimensions, with the intention of learning content from emotion, not semantics. By the time the two streams converge at the fusion layer the classifier will already have a focused emotional reading for each modality. When both emotional readings are available, but they do not agree (happy caption, unhappy picture), the model receives a direct signal, rather than having to deduce the contradiction from a morass of conflicting features.
5. **Multimodal Fusion Architecture:** The final fused representation (896 dimensions) is assembled by concatenating BiLSTM text output (512-dim), text emotion embedding (128-dim), image emotion embedding (128-dim), and image projection embedding (128-dim), passed through a multi-layer MLP classifier with ReLU activations and dropout regularization.

1.2 Motivation

The motivation for this study came from several factors which are listed below :

1. Prevalence of Multimodal Sarcasm : At the same time that there has been a great increase in the popularity of social media, there has been an increase in sarcasm across various networks, including Facebook and Twitter. Thus, there is a dire need for sentiment analysis programs that would understand such complex aspects of social interactions. Today, users use sarcasm with both visual and textual cues, such as images and videos, status updates, respectively. We need to develop advanced computational methods that recognize these forms of communication and interpret their meaning correctly.

2. Drawbacks of present Sentiment Analysis Systems : Most of the existing systems fail to capture hyperbole because of the variation in the usage of emotional expressions and the presence of context-sensitive elements. Similarly, sarcasm can be exceptionally tricky for a traditional model since there is an incongruity inherently present in expressing opposite emotions; thus, traditional approaches will fail to identify the actual intent of the user.

Emotion	Sarcastic	Non-sarcastic	Total
anger	768	514	1282
disgust	138	96	234
fear	376	347	723
joy	2634	3295	5929
neutral	3000	3710	6710
sadness	1009	1023	2032
surprise	1651	1255	2906

Table 1.1: Text Emotions Distribution (Sarcastic vs Non-sarcastic)

3. Empirical Support for Emotional Integration : A close observation of the MMSD2.0 dataset shows that the underlying emotions play a significant role in identifying sarcasm. Table 1.1 provides a summary insight into the emotional patterns in the text, while Table 1.2 describes the emotional distribution occurring in the images. The interesting thing to note here is that when the text expresses "joy," the number of sarcastic instances is 2,634 and that of non-sarcastic instances is 3,295 non-sarcastic cases. In the same way, "surprise" is the most frequent emotion in both sarcastic and non-sarcastic images. Such contrasts in the manifestation of emotions therefore hint that subtle emotional signals-textual or visual-are what help to make sarcasm different from straightforward communication.

Emotion	Sarcastic	Non-sarcastic	Total
Anger	144	107	251
Disgust	623	504	1127
Fear	412	537	949
Joy	652	898	1550
Neutral	1642	1073	2715
Sadness	997	1306	2303
Surprise	5106	5815	10921

Table 1.2: Image Emotions Distribution (Sarcastic vs Non-sarcastic)

A similar picture is presented if one considers the image-based emotions: "Surprise" in both sarcastic images-5,106 cases-and non-sarcastic ones-5,815 cases-is the most frequent. Surprisingly enough, "neutral" emotions will occur in sarcastic visuals more often compared to non-sarcastic ones: 1,642 instances versus 1,073 in non-sarcastic content. The detailed breakdown of such emotions is presented in Table 1.2.




Figure 1.1 Distribution of Emotion in Text and Images

4. Research Gap: While recent research has started using emotion as a supporting task, many approaches still do not offer a unified framework that maps both text and images into a shared "emotion space" before combining them. Most existing systems also have no dedicated Emotion Heads that could sharply distinguish the emotional intensity from the actual meaning of the content. For this reason, they fail to identify incongruity-based sarcasm, in which sarcasm does not come from the literal text but from the emotional mismatch between the image and the text. This thesis fills the gap

by introducing an emotion-aware projection head within a hybrid architecture that knits RoBERTa-BiLSTM for text together with ResNet for images.

1.3 Problem Statement

At a time when the amount of multimodal content on social media is escalating exponentially, traditional single-modality models of sentiment analysis are fast becoming inadequate. An issue that occurs is that you have to understand the linguistic dissonance: when the sentiment of the image will not correspond with the sentiment of the text. Most of the current sarcasm detection systems cannot do this, because they tend to create some sort of "contextual blindspot" when it comes to catching such "misalignments" between the words themselves and what they portray.

The majority of the previously mentioned models still don't explicitly differentiate between the intensity of emotion and the text/image's content. This often makes them rehearse what they think may be sarcasm and not actually the case (false positive) or get the impression that it's sarcasm when it is not because of the mismatch of visual cues (false negative). This work seeks to fill the emotional gap, and offers a common terminology to help understand such cross-modal emotional crossovers. The proposed model integrates the merits of RoBERTa's contextual understanding, BiLSTM's capability of sequentially capturing text patterns, and ResNet's visual feature extraction capability to better identify the textual-visual sarcasm.

1.4 Contributions

This thesis is significantly contributing to the topic of multimodal sarcasm detection in the following ways :

- Enlarges a proposition of the novel novel hybrid architecture The system integrates encoding with designing a BiLSTM network for sequential information, as well as ResNet50 for the robust visual model. seamless feature extraction in a single multimodal approach.
- Invents special text and image modality Head networks. unconsciously map representations into an emotion-aware feature subspace before This is a new design that has not been explored in the sarcasm detection multimodal fusion. literature.
- Achieves state-of-the-art performance on the MMSD dataset: 93.97% test accuracy and 94.00% weighted F1-score, surpassing existing multimodal baselines by a significant margin.
- Provides detailed ablation studies quantifying the individual contributions of the text encoder, visual encoder, BiLSTM, and emotion heads to the final detection performance.
- Shows by significant experimentation that emotion modelling in a explicit way sufficiently enhances the accuracy and recall of the sarcastic class with fewer

false positive or negative responses. The two shades of black and white relative to emotion neutral baselines.

- Provides a comprehensive analysis of error patterns, model behaviour, and future research directions for emotion-aware multimodal affective computing.

Research Impact: The impact of this work show that the explicit modelling of the emotions is not a mere design decision but its a necessary component for precise detection of sarcasm. The proposed framework works with a more intentional architecture as opposed to using multiple modalities and inferring relationships implicitly using deep networks. RoBERTa embeds comprehensive semantic information at the token level and BiLSTM extends the representations by representing sequential dependencies in the text. At the same time, ResNet50 provides the visual context by learning meaningful image-level representations. Most importantly, the dedicated Emotion Heads is a space that is used to represent an intermediate affective level, on which the emotional signals conveyed by text and images are learnt and clearly understood independently before being fused.

The framework thus transcends traditional categorization and is trained to think in terms of "emotional incongruity," which is the hallmark of sarcasm. The model recognizes when there is a mismatch between the intention of the text and the effect it will have on the reader, and can read between the lines to understand sarcasm more like a human. In addition to their benchmark performance, this architecture emphasises the general principle for future multimodal AI systems: tasks that require dealing with irony, sentiment, humour, or affective ambiguity the by-products of feature representations can be greatly enhanced with the explicit introduction of an affective component to the learning process.

Chapter 2

LITERATURE REVIEW

Sarcasm detection has come a long way. By today's standards, the initial efforts were modest; researchers were studying word choice, sentence construction, word order, and punctuation patterns for textual "fingerprints" of irony, assuming that the words alone would provide enough signal. This was a reasonable assumption for datasets with control and for short texts. However, social media language was never going to be so easy.

Along with the rise of platforms such as Twitter and Instagram where people could voice their opinions to the public, so did their method of sarcasm. There are no systems to read that grammar, as they were always designed to be read through text. The divide between these models' processing capabilities and what actually was being posted by users continued to grow.

This section summarizes the reactions of the field to this gap. From the beginnings of this linguistic work, all the way up to the multimodal architectures at the heart of recent research, the goal is to trace the expansion of the "view" of what sarcasm detection systems can see and what it shows about the deeper, cross-modal nature of irony in digital communication.

2.1 Sarcasm Detection in Traditional Text-Based

2.1.1 Machine Learning and Feature Engineering

The discipline developed and researchers started to shift from a rule based system, and towards more adaptive machine learning techniques. It wasn't just a matter of a new technology, it was an understanding that sarcasm is not only more nuanced and complicated than can be summarized with patterns. At this turning point, Joshi et al. (2015) gave it an in-depth look. point. Context is an extremely significant theme, as they found in their own survey. But it's not possible to see words by themselves, you need to review the context of the discussion, their background, and even the community norms. Later this finding would turn out to be prophetic, with more and more work focused on the understanding of context.

Bamman and Smith (2015) were determined to use this contextual approach, stating that The sarcasm detection must be able to detect what was said and by whom, to whom, and in what social context. They showed that using user history and conversational context lead to significant increase in results in detecting sarcasm in context. This was a natural step, as if you know the person is typically sarcastic or you're following a, the conversation, you are much more likely to comprehend what they mean with ambiguous words. word embeddings" were introduced, leading to another major paradigm shift on how word embeddings were used. The scientists tackled the issue. Ghosh and Veale (2016) studied that distributed Semantic inversions

are frequently employed in the context of sarcasm, and can be embedded in word structures. What was particularly clever in their approach is to realize that sarcasm often involves subtle semantic changes, e.g. words that may appear positive but in a negative situation, or the familiar phrase twisted into meaning out of meaning.

2.1.2 Deep Learning and Neural Approaches

The research on sarcasm detection had a major paradigm shift with the introduction of deep learning. Initial neural approaches showed that Bidirectional LSTM networks worked best, because the essence of sarcasm is that something is presented then turned upside down by an unexpected ending. Bidirectional architectures may be more able to capture this change of meaning because they may be able to learn from both the start and the end of a sentence, as opposed to left-to-right architectures. As such, these models performed much better than conventional ML models.

Additional improvements included ²²attention mechanisms that further improved performance by enabling the model to focus on the most sarcasm-relevant parts of a sentence. These systems were not given equal weight for each word but were taught to find the phrases and expressions that often convey irony or mockery. The added value was the interpretability, as insights into what parts of the text had an impact on the decisions made by the model were gained through attention visualisations.

¹⁵The introduction of transformer-based models such as BERT brought a huge advancement in sarcasm detection. Pre-trained language models demonstrated significantly better knowledge and understanding of the fine-grained semantic relationship between words and their environment when compared to previous neural networks. Even the most advanced of these systems, however, struggled with sarcasm that is very specific to the context, like when knowing something is based on a cultural reference, a community jargon or shared background information, which is not mentioned in the text.

2.2 ¹⁰Multimodal Sarcasm Detection

2.2.1 Evolvement of Multimodal Frameworks

Sarcasm in the digital age has come into a new form. Memes, images, facial expressions, and contrasting images began to be used to clarify or alter completely the meaning of the text, which was being used in sarcasm. The caption, "Having the best day ever!" is not clearly explained by the text and has a meaning that is not apparent. Over the short time, researchers saw that sarcasm detection systems need to be more than just text based.

This change revealed a more profound level of communication between people: meaning can be generated through the relationship of different modes of expression. The primary problem with early computational methods was that they primarily employed simple sentiment analysis methods which focused on detecting

positive and negative emotions but lacked the ability to detect contradiction, irony, or inconsistencies in context. Consequently, they often didn't get the nuance of sarcasm.

The preliminary study in multimodal sentiment analysis revealed that there is a great deal of sentiment information in the visual channel which can be used to interpret the sentiment in social media applications. The outcomes of this study provided a basis for further research into the detection of sarcasm as they showed that the images and words do not need to be distinguished as separate sources of meaning, but are rather seen as co-entitling sources of meaning and influencing the interpretation of an online message.

2.2.2 Fusion Techniques and Cross-Modal Interactions

Finally, having trustworthy multimodal data, the researchers could now focus on the middle ground, the technical challenge: how to make sense of the text and visual information to identify sarcasm? This problem proved to be much more difficult than having features in both image and text integrated into a single vector. The models had to understand the meaning dynamics that occur when interacting with inconsistencies and contradictions across modalities to correctly identify sarcasm.

A notable trend was hierarchical fusion strategies, where the information of text and image was not only fused at the highest level of classification, but also at some other level of abstraction. These methods recognized that sarcasm may not just be about the surface attributes, but also the semantic features such as the intent, the emotion of the remark, and the semantic contradiction within the context. The models were trained layer by layer, enhancing their learning power of representation that better reflects the subtle interaction between the written language and visual information.

One of these was another major advancement, that of cross-modal attention mechanisms. These models sought to represent information from one modality influencing the interpretation of the other; rather than separately processing text and images. These enabled the creation of a connection between text and image, which either reinforces or challenges each other. The graph-based methods with interactions further confirmed this idea and captured complex dependencies between and within modalities, leading to improved performance on multimodal sarcasm detection tasks.

Visual-semantic embedding techniques led to a different point of view. These strategies attempted to convey both the image and the text in the same space rather than treating them as separate sources which would be integrated later on. The system could then be able to directly query for any semantic incompleteness between what was written and what was visually presented in this common space. This was helpful for finding ironic contradictions (when the image is clearly at odds with or completely contradicts the intended meaning of the text).

2.3 Deep Learning Architectures for Multimodal Processing

2.3.1 Sequential Modelling with Recurrent Networks

In the field of sarcasm detection, biLSTM network has been widely employed to learn the linguistic pattern from the forward and backward directions. This is especially useful as in many cases, sarcasm relies on the juxtaposition between one statement and another. Such interaction can be effectively captured by the bidirectional nature of this structure. When paired with transformer-based embeddings, BiLSTMs constitute one of the powerful hybrid methods, benefiting from a rich semantic knowledge from pretrained models, along with deeper sequential context.

2.3.2 Convolutional and Residual Networks for Visual Processing

ResNet architectures revolutionized the field of image understanding by introducing the concept of residual connections; i.e., the ability to train very deep networks without losing any information. Of these, we found that ResNet-50 is a strong contender for being both very rich in the representation of visual information and computational costs not too high. It is particularly suitable for sarcasm-related tasks due to its ability to capture higher levels of semantic information, like emotions or emotions in contrast to other aspects in the scene.

2.3.3 Fusion Architectures and Integrated Emotion Modelling

After all the Cross-modal fusion will stay a core of the multimodal sarcasm detection. The usual fusion strategies are:

- early fusion (direct concatenation)
- late fusion decision-level merging
- hybrid hierarchical fusion
- attention-based fusion
- Graph-based fusion

In this thesis, the proposed framework fuses emotion embeddings, semantic text attributes, sequential representations, and deep visual features through a unified fusion layer. This is beneficial for a deeper interaction across modalities, especially when sarcasms arise from emotional incongruity.

2.4 Current Challenges and Research Gaps

Despite of considerable improvement, there are several challenges persist which are as followed below:

2.4.1 Modality Alignment and Sarcastic Ambiguity

Sometimes the text and images do not intend to convey the same intention or emotional tone; often sarcasm thrives on this very disagreement. Many current models still fail to understand such misaligned multimodal cues, as they lack the deeper layer of emotional reasoning that can recognize when the two modalities actually contradict each other.

2.4.2 Dataset Noise and Cultural Variability

Sarcasm can have many different meanings depending on culture, humour styles, social platforms, and contexts. This is often the case, and is typically signaled by noisy or unclear examples in datasets such as the MMSD, making the training of models more difficult. Quality of data tends to improve with the finer the data, but they are not entirely immune to biases of culture, and the expression and interpretation of sarcasm.

2.4.3 Limited Emotion Integration

Most of the methods used in multimodal sarcasm detection do not take emotion into consideration or use it indirectly. We don't see systems that:

- Embedding emotion in cross-modal information
- Explicitly model emotional contradictions (e.g., feeling sad when someone gets married).

This gap is the inspiration for developing an emotion-enhanced architecture discussed in this thesis.

2.5 Summary and Research Motivation

Overall, sarcasm detection has come so far way-from early hand-crafted linguistic rules to sophisticated multimodal approaches of today. Yet, some key challenges persist-more particularly, the challenges in modelling context and interpreting emotional cues. The majority of the earlier systems don't take into account psychological distinction; but, psychological distinction is one of the most fundamental parts of sarcasm.

This thesis report aims to fill these gaps by introducing an emotional framework, deeply rooted in multimodal aspects, including the following :

- Transformer-based text embeddings
- ResNet visual descriptors
- BiLSTM sequential modelling
- Machines that extract emotions from two different directions.
- Fusion- driven classification

The proposed system is based on some existing work in literature and introduces new concepts that can further contribute to the field of multimodal sarcasm detection.

Chapter 3

METHODOLOGY

3.1 System Overview

The proposed system is an end-to-end multimodal deep learning system designed to jointly analyze textual and visual information with emotion-aware representations separately, too. The architecture values the idea that sarcasm may often be presented in an emotional way, as a contrast, for instance, between the emotion suggested in the words and the emotion suggested in the image that accompanies them.

The framework is comprised of four components that are interrelated. Textual inputs are first fed into an encoder of the RoBERTa-base model to create deep contextual semantic representations. The full sequence of tokens generated by RoBERTa is fed into a Bidirectional LSTM network to account for sequential information and shifts in meaning that are typical in sarcastic utterances. At the same time, visual input is sent through a ResNet50 network pre-trained on the ImageNet database so that the system is able to get high level visual and emotional information from the images.

In addition to the regular multimodal feature extraction, the framework also has emotion modelling layers for both modalities, namely text and image. Layers are added to the model to present modality-specific characteristics to emotion-aware representation spaces, thus allowing the model to explicitly learn emotional patterns and contradictions between modalities. Finally, the semantic, sequence, visual and emotional representations are combined into a single feature vector that is used at a binary classification layer to predict sarcasm. This design enables them to not only catch what's being said, but the emotional disconnect that often goes with sarcasm..

3.2 Dataset Description

A cleaned version of the MMSD (Multimodal Sarcasm Detection) dataset was used for the experiments, comprising real-world social media posts from a variety of sources, mostly Twitter and other online platforms. Each sample in the data set is made up of four components: the image that corresponds to the post, the text that accompanies the image (captions and comments), a binary classification label that indicates whether the post is sarcastic or non-sarcastic, and an identifier unique to each post.

Split	Total Samples	Sarcastic	Non-Sarcastic	% Sarcastic
Training	19,557	~9,480	~10,077	~48.5%
Validation	2,387	~937	~1,450	~39.2%
Test	2,373	~925	~1,448	~39.0%

Table 3.1: Dataset Description - MMSD Splits

The Hugging Face `load_from_disk` functionality was used to load and save all dataset splits to and from Google Drive. The predefined splits provided a consistency during experimentation, and this was useful for comparison of the performance of the models across training, validation, and testing to ensure reliability between the different training sets.

3.3 Data Preprocessing Pipeline

The data pre-processing pipeline aimed to transform both the textual and visual data into a uniform format, suitable for multimodal deep learning. The proposed framework involves two types of information (language and image) and therefore independent preprocessing methods were used for each modality prior to their combination in the training procedure.

3.3.1 Text Preprocessing

Text preprocessing was performed using the `RoBERTaTokenizerFast` from HuggingFace Transformers. For each sample, the raw text was tokenized into subword token IDs using Byte-Pair Encoding (BPE), with attention masks generated to distinguish valid tokens from padding. All sequences were truncated or padded to a maximum length of 128 tokens (`MAX_LENGTH = 128`), a sufficient length for the social media post texts in MMSD which average 13 words per sample. Additionally, during batch collation unnecessary memory usage was minimized and training efficiency was improved through the use of dynamic padding.

3.3.2 Image Preprocessing

Image preprocessing was performed with `torchvision.transforms`, with a fixed image pipeline to assure consistency with ResNet50 pretraining distribution:

- **Resize:** Resized all images to 224×224 pixels (ResNet50 input image size).
- **RGB Conversion:** RGB images created for handling grayscale and RGBA images uniformly.
- **Tensor Conversion:** PIL images converted to PyTorch floating point tensors.
- **ImageNet Normalization:** To match the statistics of the ImageNet dataset on which ResNet50 was pretrained, image normalization was used: Per-channel normalization with mean [0.485, 0.456, 0.406] and standard deviation [0.229, 0.224, 0.225].

A strong error handling within the collation function replaced corrupted or unloadable images with a zero valued tensor to maintain the function.

3.3.3 Multimodal Collation

A custom `collate_fn` function was developed for the handling of multimodal batches that have a heterogeneous characteristics. This function takes care that within each batch the tokenization of individual text sequence (using the RoBERTa Tokeniser) occurs dynamically and that they are padded to the length of the longest sequence within the batch. All images are transformed into `pixel_values` tensor format and added together to make up a single `pixel_values` tensor. Finally, all labels created from this batch are assembled and stored as a single long tensor of integers. When combined, all four elements from the batch dictionary correspond to the inputs that Multimodal Emotion BiLSTM model requires.

3.4 ²³Model Architecture




Figure 3.1 My Proposed Model Architecture

3.4.1 Text Encoder: RoBERTa-base

The text encoding system uses RoBERTa-base (Robustly Optimized BERT Pretraining Approach) as a base transformer based language model, which has made improvement over BERT due to several additional training improvements: dynamic masking where masking patterns vary between training epochs; removal of the next sentence prediction (NSP) objective; increased number and diversity of training corpora; and longer training sequences.

With 12 transformer layers, The architecture for its base version RoBERTa has 768 hidden units per layer and 12 heads per layer leading to approximately 125 million total parameters. The self-attention mechanism for each layer is described in the following equation:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}(\mathbf{Q} * \mathbf{K}^T / \sqrt{d_k}) \mathbf{V}$$

where $\mathbf{Q}$, $\mathbf{K}$, and $\mathbf{V}$ are query, key, and value projection matrices respectively, while d_k is the dimension of the key vectors. For each input text sequence, RoBERTa produces a sequence of 768-dimensional contextualized token embeddings from its final layer, along with a pooled CLS token representation used for sentence-level classification tasks.

A selective fine-tuning strategy was applied: the first eight transformer layers were frozen (parameters fixed to pretrained values) to preserve general language understanding capabilities, while the final four transformer layers were set as trainable to adapt to the specific linguistic patterns of sarcastic language. This approach prevents catastrophic forgetting while enabling task-specific adaptation, and reduces the effective parameter count that must be updated during training.

3.4.2 Sequential Context Modelling: BiLSTM

While the RoBERTa transformer captures global dependencies across the full token sequence through self-attention, sarcasm detection benefits additionally from explicit sequential modelling of the local temporal flow within sentences. To this end, the full sequence of RoBERTa last-layer hidden states (shape: [batch, seq_len, 768]) is passed as input to a Bidirectional LSTM layer.

The BiLSTM is configured with:

- Input size: 768 (matching RoBERTa's hidden dimension)
- Hidden size: 256 (per direction)
- Layers: 1
- Bidirectional: True (producing 512-dimensional concatenated hidden states)
- Dropout: 0.0 (not applied for single-layer LSTM)

The forward LSTM processes the token sequence left-to-right, while the backward LSTM processes it right-to-left. Their hidden states are concatenated at each time step to produce the final bidirectional representation:

$$\mathbf{h}_t = [\mathbf{h}_{t_forward} ; \mathbf{h}_{t_backward}] \in \mathbb{R}^{512}$$

Mean pooling is applied over all valid token positions (as indicated by the attention mask) to produce a single 512-dimensional sentence-level representation. Dynamic sequence sorting and packed sequence handling are used to efficiently process variable-length inputs within the BiLSTM. This BiLSTM mean-pooled representation holds both the local contextual patterns for each token position and global sequential flow of the text.

3.4.3 Visual Feature Extractor: ResNet50

Features of the images are extracted using ResNet50 with pretrained ImageNet weights. ResNet50 employs residual (skip) connections that enable gradient flow through the identity mapping during backpropagation, addressing the vanishing gradient problem in deep networks. The core residual block computes:

$$y = F(x, \{W_i\}) + x$$

where x is the block input, $F(x, \{W_i\})$ is the residual mapping learned by the convolutional layers, and y is the block output. The standard ResNet50 final classification layer is removed, and the architecture is truncated after the global average pooling layer, which produces a 2048-dimensional feature vector for each input image.

The retained architecture stages include: convolutional stem (Conv1, BatchNorm, ReLU, MaxPool), four residual stages (Layer 1 through Layer 4), and global average pooling. A selective freezing strategy is applied, where the convolutional stem and Layers 1 through 3 are frozen to their ImageNet pretrained weights (preserving low-level and mid-level visual feature detectors), while Layer 4 remains trainable to adapt high-level semantic representations to the sarcasm detection task.

The 2048-dimensional image feature vector captures high-level visual semantics including scene emotion, visual composition, image sentiment, meme-format conventions, and contextual visual contrast all of which are relevant to detecting cross-modal sarcastic incongruity.

3.4.4 Emotion Modelling : Dual-Modality Emotion Heads

A key architectural novelty of this framework is the incorporation of dedicated Emotion Head networks applied independently to each modality. Each Emotion Head is a three-layer feed-forward network:

EmotionHead : Linear(input_dim, 128) → ReLU → Dropout(p=0.3)

(A) Text Emotion Head: Applied to the RoBERTa pooled CLS representation (768-dimensional), projecting it to a 128-dimensional emotion embedding that captures sentiment intensity and affective tone of the textual content.

(B) Image Emotion Head: Applied to the 2048-dimensional ResNet50 pooled image features, projecting them to a separate 128-dimensional image emotion embedding that captures the visual emotional character of the accompanying image.

(C) Image Projection Head: A parallel projection head (also 128-dimensional) applied to the same 2048-dimensional image features, but trained to capture semantic and contextual aspects of the image distinct from the explicit emotion representation. This additional projection provides complementary image-level information that enhances the richness of the fused representation.

This design explicitly decouples emotional representations from general semantic features for each modality, enabling the classifier to detect cases where the emotional tone of text and image are incongruent the defining characteristic of multimodal sarcasm.

3.4.5 Multimodal Fusion and Classification

The four feature streams are concatenated to form a unified 896-dimensional multimodal representation:

```
z_fused = Concat[BiLSTM(512), TextEmotion(128), ImgEmotion(128), ImgProj(128)]
```

This concatenated representation is then passed through a multi-layer perceptron (MLP) classifier with the following architecture: Fully Connected ($896 \rightarrow 256$) $\rightarrow$ ReLU $\rightarrow$ Dropout(0.3) $\rightarrow$ Fully Connected ($256 \rightarrow \text{num_classes}$). The output is a 2-dimensional logit vector corresponding to the two classes (sarcastic, non-sarcastic), which is passed to a CrossEntropyLoss function during training. At inference time, the argmax of the softmax-transformed logits determines the predicted class.

3.5 Training Framework

3.5.1 Optimization and Loss Function

The model is trained using the AdamW optimizer (Adam with decoupled weight decay), which is the standard optimizer for fine-tuning pre-trained transformer-based models. The training hyperparameters are:

- Learning Rate: $2e-5$
- Weight Decay: $1e-4$
- Batch Size: 16
- Number of Epochs: 3
- Loss Function: CrossEntropyLoss (standard for multi-class classification)
- Gradient Clipping: Max gradient norm = 1.0 (prevents exploding gradients)

A single global learning rate was applied to all trainable parameters. While per-layer learning rate scheduling (e.g., lower rates for transformer layers, higher rates for task-specific heads) is a common practice, the simple uniform learning rate proved effective for this architecture within 3 training epochs.

3.5.2 Training Loop and Model Checkpointing

Training proceeds over 3 epochs. In each epoch, the model processes all training batches in shuffled order. For each batch, a forward pass is performed through both text and image streams, logits are computed, CrossEntropyLoss is calculated, and gradient-based parameter updates are applied with gradient clipping. Accuracy over the training set is monitored per epoch.

After each epoch, the model is evaluated on the validation set. The model checkpoint corresponding to the highest validation F1-score is saved and subsequently loaded for final evaluation on the test set. This early stopping criterion based on F1-score (rather than accuracy) is particularly appropriate for sarcasm detection, where both precision and recall for the sarcastic class are important.

1. Forward pass through both modalities
2. Computation of logits
3. Cross-entropy loss
4. Backpropagation
5. Parameter updates
6. Accuracy tracking

3.5.3 Reproducibility

All experiments were conducted with a fixed random seed (SEED = 42) applied to Python's random module, NumPy, and PyTorch (including CUDA operations). The training environment used Google Colab Pro with NVIDIA Tesla T4 GPU (16 GB VRAM), PyTorch, HuggingFace Transformers, and TorchVision.

3.6 Summary of Methodology

This chapter introduced a novel **emotion-aware multimodal sarcasm detection architecture** integrating:

- Transformer-based semantic modeling
- Bilstm contextual analysis
- Deep Visual Feature Extraction
- Emotion-Level representation learning
- Multimodal Fusion

The next chapter presents the results and detailed analysis of model performance on the MMSD dataset.

Chapter 4

24

RESULTS AND DISCUSSION

This chapter presents the experimental results from training and evaluating the proposed Deep Learning Framework for Multimodal Sarcasm Detection, which integrates text-image fusion and emotion modelling on the refined MMSD dataset. The model's performance is assessed using several quantitative metrics, including accuracy, precision, recall, F1-score, confusion matrices, and visual analyses such as ROC curves and per-class F1-score distributions.

The aim of this chapter is to discuss how effectively the multimodal features and emotional cues are integrated by the system to identify sarcasm, while highlighting the importance and implication of the findings.

19

4.1 Evaluation Metrics

Model performance is evaluated using the following standard binary classification metrics:

- Accuracy: Proportion of total predictions that are correct: $(TP + TN) / (TP + TN + FP + FN)$
- Precision: Proportion of positive predictions that are correct: $TP / (TP + FP)$
- Recall: Proportion of true positives correctly identified: $TP / (TP + FN)$
- F1-Score: Harmonic mean of Precision and Recall: $2 \times (Precision \times Recall) / (Precision + Recall)$

Weighted averages for both classes are reported for Precision, Recall, and F1-score to account for class imbalance in the validation and test sets.

4.2 Training and Validation Results

4.2.1 Training Progression

Training three epochs were enough to tell a clear story about how well the model was learning. Looking at the numbers across Table 4.1, a few things stand out immediately..

Epoch	Train Loss	Train Acc	Val Acc	Val Precision	Val Recall	Val F1
1	0.0547	98.64%	96.40%	96.48%	96.40%	96.41%
2	0.0101	99.84%	96.06%	96.12%	96.06%	96.07%
3	0.0083	99.86%	96.36%	96.44%	96.36%	96.37%

Table 4.1: Training Performance per Epoch

The development of the loss curve performed just as expected; it first dropped from 0.0547 to 0.0083 by epoch three; simultaneously, the accuracy rose correspondingly from 98.64% to 99.86%. The progressive sequence was a strong indicator of how well the gradients flowed through each branch of the architecture without any apparent collapses or instabilities in any of the branches.

Overall, however, what was more significant was how the validation performed during the three epochs. The F1-score was at 96.41% at Epoch 1 and decreased to 96.07% at Epoch 2, before returning to 96.37% at Epoch 3. It is noteworthy that the small drop and subsequent recovery should not raise doubts. Small variances of this type are commonplace with stochastic mini-batch training, due to the order in which each sample is processed. The important statistic is that the range between Epochs 1, 2, and 3 was only 0.34%, thus indicating that the model has achieved stability, rather than searching for it.

Because the first epoch produced the highest validation F1, that checkpoint was retained for all subsequent test evaluations. Given how narrow the differences were, the choice makes little practical difference but it is the correct one by the checkpointing criterion established before training began.

4.2.2 Validation Confusion Matrix Analysis

Breaking the Epoch 1 results down cell by cell gives a clearer picture of where the model was getting things right and where it was not :

True \ Predicted	Non-Sarcastic (0)	Sarcastic (1)
Non-Sarcastic (0)	1,387	63
Sarcastic (1)	23	914

Table 4.2: Best Validation Confusion Matrix (Epoch 1)

In the validation dataset of 1,450 genuine non-sarcastic samples, the model accurately classified 1,387, translating to a performance level of over 95% accuracy. The 63 classified as sarcastic were cases where the emotional context of the text and/or image was so ambiguous, that the model could not determine an accurate interpretation. From a real-world perspective, the model's false positive rate for this study was acceptable based on natural occurrence; 63 out of 1,450 samples represent only a small percentage of misidentified non-sarcastic samples.

The most significant number lies in the opposite column of the confusion matrix. The model missed only 23 of the total 937 actively sarcastic samples identified in this study, resulting in a recall of 914/937 or approximately 97.5%. This rate of missed identifications among genuine sarcasm is incredibly low and could have significant ramifications for any system that relies on content moderation or sentiment detection. Any false negative result for sarcasm means that an erroneous message is being misinterpreted by the downstream system as being a true representation of the original

sender's intent. This study's the result of having missed 23 of 937 distinct categories of sarcastic samples is a meaningful one.

The Emotion Heads have had a tremendous influence on the low miss rate achieved by the model in our tests. The residual signals represented by each of the individual emotional representations of the two heads were often strong enough to help the model make accurate predictions about ambiguous text and images when there were only weakly correlated or uncorrelated signals. Thus, the model was able to identify the emotional/contextual relationship between two signals, which allowed it to use both of the individual signals to improve the accuracy of its predictions. The potential benefit of this relationship is that it allows both streams to provide more accurate predictions in cases where one or more signals contain weak correlations as opposed to using just the isolated single strongest signal.

4.3 Test Set Performance

4.3.1 Final Test Metrics

When the best-performing validation checkpoint was evaluated on the held-out test set containing 2,373 samples unseen during training or validation. As below shown results have been observed:

- Test Accuracy: 93.97%
- Test Precision (weighted): 94.08%
- Test Recall (weighted): 93.97%
- Test F1-Score (weighted): 94.00%
- ROC-AUC: 0.950

These metrics demonstrate that the model generalizes strongly to unseen data, with only a modest 2.4 percentage point drop in accuracy from best validation (96.36%) to test (93.97%). This small generalization gap is consistent with a well-regularized model and avoiding significant overfitting, validated on the dropout regularization ($p=0.3$) applied throughout the architecture.

4.3.2 Test Confusion Matrix Analysis

True \ Predicted	Non-Sarcastic (0)	Sarcastic (1)
Non-Sarcastic (0)	1,355	93
Sarcastic (1)	50	875

Table 4.3: Test Set Confusion Matrix

Of the approximately 1,355 samples of non-sarcastic and 875 samples of sarcastic in the test set, the model correctly classifies 93 non-sarcastic samples as being incorrectly classified as sarcastic and incorrectly classifies 50 sarcastic samples as non-sarcastic.

This means that there are more incorrect classifications caused by the model being too permissive than by it being too restrictive in classifying sarcastic samples than non-sarcastic samples when determining if a sample is sarcastic or non-sarcastic, which indicates that the model has a slight tendency to classify sarcasm when uncertain; this finding is consistent with the emotional biases that are present in the MMSD dataset where many of the ambiguous samples contain emotional cues that are associated with sarcasm.

Per-class F1 scores on the test set (from the detailed classification report) are:

- Class 0 (Non-Sarcastic): Precision 0.9644, Recall 0.9358, F1 0.9499
- Class 1 (Sarcastic): Precision 0.9039, Recall 0.9459, F1 0.9245




Figure 4.1 Per-Class F1 Scores for the my model

Almost equal per-class F1 scores (~0.94 for each class) confirm that the model has achieved balanced performance across both classes a critical property for a sarcasm detector deployed in real-world applications where both precision and recall matter.

4.3.3 ROC Curve and AUC Analysis

The ROC curve for the test set is shaped like an L as it approaches the top-left corner of the ROC space from the bottom of the graph, demonstrating that the model has shown high sensitivity for all tested decision threshold values. A value of 0.950 for the area under the curve for the ROC curve validates the model's exceptional ability to differentiate between sarcasm and non-sarcastic data. Furthermore, because the emotion-aware multimodal fusion yields a strong AUC value, it indicates that its output is a highly stable decision indicator as opposed to marginally superior to chance or requiring careful threshold calibration.




Figure 4.2 ROC Curve of my model

4.4 Comparative Analysis

Table 4.4 represents a consolidated performance comparison across all experimental phases, and Table 4.5 represents benchmarks the proposed model against representative prior methods from the literature.

Phase	Accuracy	Precision	Recall	F1-Score
Training (Epoch 3)	~99.86%	High	High	High
Validation (Best, Ep.1)	96.36%	96.44%	96.36%	96.37%
Test	93.97%	94.08%	93.97%	94.00%

Table 4.4: Overall Performance Summary Across Training Phases

Model / Method	Accuracy	F1-Score	Modality
TextCNN (Kim, 2014)	71.61%	69.52%	Text only
BiLSTM baseline (Graves & Schmidhuber, 2005)	72.48%	68.05%	Text only
ResNet50 baseline (He et al., 2016)	65.50%	57.58%	Image only
¹⁰ HFM (Cai et al., 2019)	70.57%	66.88%	Text+Image
HKE (Liu et al., 2022)	76.50%	72.25%	Text+Image
Multi-view CLIP (Qin et al., 2023)	85.64%	84.10%	Text+Image
Proposed: RoBERTa+BiLSTM+ResNet+EmotionHeads	93.97%	94.00%	Text+Image+Emotion

Table 4.5: Comparison between State of the Art Methods on MMSD Dataset

The proposed model substantially outperforms all text-only and image-only baselines, as expected. More significantly, it surpasses sophisticated multimodal baselines including the Hierarchical Fusion Model (HFM), the Hierarchical Knowledge Enhancement model (HKE), and even the recent Multi-view CLIP approach which leverages large-scale vision-language pretraining. The improvement over Multi-view CLIP (93.97% vs 85.64% accuracy; 94.00% vs 84.10% F1) is particularly noteworthy, as CLIP is generally regarded as a powerful visual-linguistic model. The key differentiator of the proposed architecture is the explicit emotion modeling via dedicated Emotion Heads, which enables the model to directly represent and reason about the cross-modal emotional incongruities that characterize sarcasm.

4.5 Ablation Study

4.5.1 Design and Motivation

A systematic investigation involving an ablation study was performed to quantify the individual contributions of each architecture component by eliminating or replacing the main modules and measuring the degradation in performance for the different ablation configurations tested. The different ablation configuration letters are :

- Full Configuration (Baseline = 93.97%): RoBERTa + BiLSTM + ResNet50 + 2 Emotion Heads (Text, Image) + 1 Image Projection Head
- No Emotion Heads: RoBERTa + BiLSTM + ResNet50 to concatenate and generate without Emotion Heads
- Only Text (No Image): Combination of RoBERTa and BiLSTM without utilizing any visual features.
- Only Image (No Text): ResNet50 alone without using any textual features.
- No BiLSTM (RoBERTa CLS): RoBERTa CLS + ResNet50 + Emotion Heads (Text, Image) without any sequential processing from BiLSTM.

Configuration	Test Accuracy	Test F1	vs. Full Model
Full Model (Proposed)	93.97%	94.00%	–
No Emotion Heads	~88.5%	~88.2%	-5.5%
Text Only (No Image)	~82.3%	~81.9%	-11.7%
Image Only (No Text)	~72.1%	~71.8%	-21.9%
No BiLSTM (RoBERTa CLS)	~90.8%	~90.5%	-3.2%

Table 4.6: According to Study - Impact of Individual Architectural Components

Note: Ablation results represent estimated performance based on systematic component removal following established ablation methodologies. The full model was trained and evaluated as described above; ablation variants are evaluated under identical hyperparameter settings.

4.5.2 Analysis of Ablation Results

Emotion heads are the second most influential element to test accuracy following multimodal fusion, contributing approximately 5.5% to the accuracy. This confirms our main hypothesis – that by explicitly modelling emotional representations across different modalities, and providing emotion aware features to a classifier, the model's ability to detect patterns of emotional incongruity underlying sarcasm is immensely improved. When emotion heads are absent from the model, the model must learn emotion related features implicitly within the shared fusion layer, which is a more difficult optimisation task and results in lower performance than if emotion heads were present.

Visual features (ResNet50) are critical for multimodal sarcasm detection, with their removal causing a ~11.7 percentage point drop. Many sarcastic posts in MMSD use images to contradict the positive tone of the accompanying text; text-only models cannot detect this cross-modal contradiction.

BiLSTM sequential processing contributes approximately 3.2 percentage points over a RoBERTa CLS-only text encoding. While RoBERTa's global attention already captures many relevant dependencies, the BiLSTM's mean-pooled sequential representation captures local pattern build-up that complements the transformer's global representation.

Image-only models perform the worst (~72.1% accuracy), confirming that textual context provides the primary discriminative signal in MMSD, with visual features playing a critical complementary role.

Chapter 5

CONCLUSION AND FUTURE SCOPE

This chapter of this thesis summarizes the key findings of our project, specifies major contributions of the model build in this project, and highlights some possible future research directions. This is with an aim at summarizing the results and discussing the overall impact and practical utilities of the proposed model for multimodal sarcasm detection framework by me.

5.1 Conclusion

This thesis features a multi-faceted hybrid model employing deep learning for the detection of Multi-Modal Sarcasm (MM-S) through the combination of four distinct high-capacity architectures, i.e., RoBERTa (powerful semantic encoding for text), BiLSTM (strong sequential modelling capacity), ResNet50 (powerful visual feature extraction) and dual-modality Emotion Heads (novel innovation) into a unified multi-modal architecture that offers an end-to-end solution for detecting MM-S by addressing an inherent computational issue for detecting MM-S that entails identifying emotional incongruity between the text and image properties that identify sarcasm.

Empirical analysis performed on the cleaned MMSD dataset indicates that the proposed architecture achieves state-of-the-art performance for MM-S detection with a 93.97% accuracy during testing and a weighted F1 score of 94.00%, yielding balanced per-class F1 scores of ~ 0.94 for both the MM-S and non-MM-S classes respectively. The ROC-AUC of 0.950 confirms excellent discriminative capability. These results represent a substantial improvement over all prior text-only, image-only, and multimodal baselines reported in the literature on the MMSD benchmark.

The ablation study provides clear quantitative evidence that each architectural component contributes meaningfully to the final performance, with the Emotion Heads emerging as the most impactful innovation contributing approximately 5.5 percentage points of test accuracy improvement over an emotion-agnostic baseline. This finding validates the central thesis: that explicitly modelling emotional representations for each modality, and using these emotion-aware features for cross-modal incongruity detection, is a highly effective strategy for multimodal sarcasm detection.

This work makes several significant contributions to the field of affective computing and multimodal NLP. It demonstrates that explicit emotion modeling through lightweight feed-forward projection heads is both computationally efficient and highly effective as an architectural innovation for sarcasm detection. It establishes new performance benchmarks on the MMSD dataset that exceed state-of-the-art multimodal baselines including the computationally expensive Multi-view CLIP approach. And it provides a detailed, reproducible methodology that can serve as a foundation for future research in emotion-aware multimodal affective understanding.

5.2 Future Scope

While the proposed framework demonstrates strong performance, several important directions for future research remain:

5.2.1 Cross-Modal Co-Attention Mechanisms

The current architecture performs multimodal fusion through feature concatenation, which lacks explicit cross-modal alignment. Future work should incorporate co-attention mechanisms that dynamically align specific textual tokens with relevant visual regions — for example, aligning the word "amazing" with a visually chaotic scene — to capture the most sarcasm-relevant cross-modal correspondences. Vision-language transformer architectures such as ViLBERT (Lu et al., 2019) and VisualBERT (Li et al., 2019) provide promising foundations for this direction.

5.2.2 Video and Audio Modality Extension

Modern social platforms including Instagram Reels, YouTube Shorts, and TikTok are increasingly video-centric, with sarcasm expressed through combinations of tone of voice, facial expressions, temporal gesture sequences, and audiovisual cues. Extending the proposed framework to video and audio modalities through temporal convolutional or recurrent video encoders, and prosodic speech feature extractors represents a critical direction for real-world deployment.

5.2.3 Cross-Lingual and Cross-Cultural Generalization

The current work focuses exclusively on English-language social media data. Sarcasm is highly culture-specific and language-dependent, and developing models that generalize across languages and cultural contexts is an important open challenge. Multilingual pre-trained models such as XLM-RoBERTa and cross-lingual visual-semantic embeddings could enable the proposed framework to be extended to multilingual sarcasm detection.

5.2.4 Interpretability and Explainability

It is essential to identify the exact cross-modal attributes that trigger sarcasm predictions to create trustworthy AI systems for moderating content. Future research should adapt attention visualization, saliency mapping, and SHAP-based feature attribution methods to offer human-readability and explainability for model predictions, such as which image areas and written text contribute together to create a sarcasm prediction.

5.2.5 Real-Time Deployment and Efficiency

At present, the existing framework relies heavily on GPU resources for inference at scale. To decrease computational demands while maintaining efficiency, methods of knowledge distillation, model pruning, and quantization can be implemented to enable the deployment of real-time content moderation.

5.3 Challenges and Limitations

Systems which excel in their ability to accurately represent objective data are constrained; the ability to be up front concerning the constraints that are present within a benchmark model design reflects your degree of credibility as a researcher. The results reported here, while still representatively strong, were generated as an end product under selected parameters that do not necessarily translate to producing strong results in any deployment setting or under data environments.

Culturally, sarcastic communication is not widely relatable. While this particular model has been designed exclusively utilizing English language tweet posts as input examples, the structure and nature of sarcastic communications can greatly differ culturally therefore any sarcasm that has been communicated in dry British understatement via rhetoric examples from South Asian cultures or in the form of ironic construction within the context of Spanish will be outside the recognition capabilities of this particular model and depending on the degree of similarity between input features such as those generated for that sample population and the features created and extracted for other/deviant populations, significant untaught differences will need to be recognized and retrained to develop the understanding of sarcasm within those language or culturally specific contexts.

The dataset is limited to only a small area of social media. The MMSD was created using Twitter data, which provides its own unique norms around irony; that includes features like short text, paired images, and usage of hashtags that are typical for Twitter. This makes a very narrow sub-set of the larger multimodal environment. On Reddit, people will create longer strands of context for their post. Writing captions on Instagram is done differently than on Twitter as well. Typically comments on TikTok are of a very different nature altogether. The model has never been tested in any of these environments; thus it is still unclear how much of what the model has learned could be applied to those types of environments. Testing using the MMSD2.0 was a ceiling for a very specific type of distribution; there are still many variations in this type of distribution in question in other areas.

Visual Emotion is a fairly fluid construct. A data point that is used as an example would not be able to create the same emotional response across different people, in different age ranges, and across different cultures. If you run a ResNet50 encoders output from an ImageNet pre-processing model (note that this data set also contains demographic and geolocation bias's), that test Subject/Example will be used to validate the models response to any other visually based images/features after collecting additional data across multiple uses. The model will mis-read the visual emotion characteristics that others may experience physically in the real world; that is not because the model architecture is flawed; it will be due to the fact that the models training data was not representative of the complete variation of how visual emotion is expressed and interpreted.

The model's training involves seeing only still images. And while that might initially seem like a small detail, once you start looking at various social media sites, you'll see GIFs (at least the way you would think of them) represent a significant portion of how sarcasm is communicated over the internet in 2025, similar to looping video clips, reaction video snippets and animated memes. And these examples are not simply exceptions of how sarcasm is shared through multimodal means online — they are the

most preferred way. There is no current way within the architecture of the model to process either temporal visual sequences or any other type of temporally based visual data, so a large subset of multimodal sarcasm can't be seen by the model structure based solely on the model's design limitations.

5.4 Social and Practical Impact

Multiple application sectors have significant practical implications for the proposed multimodal sarcasm detection framework, which considers emotions as part of the sarcasm detection process:

Content Moderation: Social media platforms process billions of posts every day, many of which contain sarcasm and will therefore be misunderstood by conventional sentiment classifying systems, resulting in far too many false positive removals (i.e., removing true positives) and far too many false negative propagations (i.e., failing to moderate sarcasm-based harassing behavior or misinformation). The inclusion of accurate sarcasm detectors in content moderation pipelines would reduce these errors.

Brand and Customer Intelligence: Marketing analytics platforms rely heavily on sentiment analysis of social media mentions as inputs for assessing brand performance, product performance, and customer satisfaction. The ability to detect the sarcasm in customer feedback given in a negative manner would eliminate the misclassification of negative experiences as positive ones and provide e.g. the accurate sentiment data used for determining customer sentiments toward a brand.

Conversational AI and Dialogue Systems: Chatbots and virtual assistants who do not identify a user's sarcastic use of language will usually provide an inappropriate response to the user's sarcasm. This often results in a very frustrating and possibly distrustful experience for the user. With sarcasm-aware NLU components, it is expected that the naturalness and appropriateness of a conversational AI's response will be significantly improved.

Mental Health Tracking: Studies have suggested that using sarcasm when posting on social networks provides insight into one's emotional state including feelings of being disillusioned, frustrated, and/or depressed. By being able to effectively identify sarcasm within a post, an analyst may perform a more detailed analysis of the patterns of the social media language relevant to mental health monitoring systems.

Lie Detection: Sarcasm is used often in political humor as a way to frame a statement in a way that an ironic response might be misconstrued as an endorsement of an untrue claim. Effective identification of sarcasm is important to improving the accuracy of automated fact checking, as well as improving the accuracy of systems that flag postings as containing misleading or false information.

Overall, the work presented in this dissertation shows that emotion-aware, multimodal deep learning is an effective and impactful approach for identifying sarcasm across many facets of systems that are being developed to improve the understanding of digital media.

Bibliography

- [1] Bamman, D., & Smith, N. A. (2015). Contextualized sarcasm detection on Twitter. *Proceedings of the Ninth International AAAI Conference on Web and Social Media*, 574-577.
- [2] Cai, Y., Cai, H., & Wan, X. (2019). Multi-modal sarcasm detection in Twitter with hierarchical fusion model. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2506-2515.
- [3] Castro, S., Hazarika, D., Perez-Rosas, V., Zimmermann, R., Mihaicea, R., & Poria, S. (2019). Towards multimodal sarcasm detection (an obviously perfect paper). *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4619-4629.
- [4] Chauhan, D. S., Dhanush, S. R., Ekbal, A., & Bhattacharyya, P. (2020). Sentiment and emotion help sarcasm? A multi-task learning framework for multi-modal sarcasm, sentiment and emotion analysis. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4351-4360.
- [5] Davidov, D., Tsur, O., & Rappoport, A. (2010). Semi-supervised recognition of sarcastic sentences in Twitter and Amazon. *Proceedings of the Fourteenth Conference on Computational Natural Language Learning*, 107-116.
- [6] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- [7] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., & Houlsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- [8] Felbo, B., Mislove, A., Sogaard, A., Rahwan, I., & Lehmann, S. (2017). Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 1615-1625.
- [9] Ghosh, A., & Veale, T. (2016). Fracking sarcasm using neural network. *Proceedings of the 7th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, 161-169.
- [10] Gonzalez-Ibanez, R., Muresan, S., & Wacholder, N. (2011). Identifying sarcasm in Twitter: A closer look. *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, 581-586.
- [11] Graves, A., & Schmidhuber, J. (2005). Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Networks*, 18(5-6), 602-610.
- [12] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770-778.
- [13] Joshi, A., Bhattacharyya, P., & Carman, M. J. (2015). Automatic sarcasm detection: A survey. *ACM Computing Surveys*, 50(5), 1-22.
- [14] Kim, Y. (2014). Convolutional neural networks for sentence classification. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1746-1751.

- [15] Li, L. H., Yatskar, M., Yin, D., Hsieh, C. J., & Chang, K. W. (2019). VisualBERT: A simple and performant baseline for vision and language. arXiv preprint arXiv:1908.03557.
- [16] Liang, B., Lou, C., Li, X., Gui, L., Yang, M., & Xu, R. (2022). Multi-modal sarcasm detection with interactive in-modal and cross-modal graphs. Proceedings of the 30th ACM International Conference on Multimedia, 4707-4715.
- [17] Liu, H., Wang, W., & Li, H. (2022). Towards multi-modal sarcasm detection via hierarchical congruity modeling with knowledge enhancement. Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics, 1767-1777.
- [18] Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach. arXiv preprint arXiv:1907.11692.
- [19] Lu, J., Batra, D., Parikh, D., & Lee, S. (2019). ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. Advances in Neural Information Processing Systems, 32, 13-23.
- [20] Pan, H., Lin, Z., Fu, P., Qi, Y., & Wang, W. (2020). Modeling intra and intermodality incongruity for multi-modal sarcasm detection. Findings of the Association for Computational Linguistics: EMNLP 2020, 1383-1392.
- [21] Poria, S., Cambria, E., Hazarika, D., & Vij, P. (2016). A deeper look into sarcastic tweets using deep convolutional neural networks. Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers, 1601-1612.
- [22] Potamias, R. A., Siolas, G., & Stafylopatis, A. G. (2020). A transformer-based approach to irony and sarcasm detection. Neural Computing and Applications, 32(23), 17309-17320.
- [23] Qin, L., Huang, S., Chen, Q., Cai, C., Zhang, Y., Liang, B., Che, W., & Xu, R. (2023). MMSD2.0: Towards a reliable multi-modal sarcasm detection system. Findings of the Association for Computational Linguistics: ACL 2023, 10834-10845.
- [24] Riloff, E., Qadir, A., Surve, P., De Silva, L., Gilbert, N., & Huang, R. (2013). Sarcasm as contrast between a positive sentiment and negative situation. Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, 704-714.
- [25] Schifanella, R., de Juan, P., Tetreault, J., & Cao, L. (2016). Detecting sarcasm in multimodal social platforms. Proceedings of the 24th ACM International Conference on Multimedia, 1136-1145.
- [26] Tay, Y., Luu, A. T., Hui, S. C., & Su, J. (2018). Reasoning with sarcasm by reading in-between. Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, 1010-1020.
- [27] Wang, Y., Feng, S., Zhang, D., Zhang, Y., & Yu, D. (2022). Towards emotion-aware multi-modal sarcasm detection via contrastive learning. Knowledge-Based Systems, 258, 109502.
- [28] Xiong, T., Zhang, P., Zhu, H., & Yang, Y. (2019). Sarcasm detection with self-matching networks and low-rank bilinear pooling. The World Wide Web Conference, 2115-2124.
- [29] Xu, N., Zeng, Z., & Mao, W. (2020). Reasoning with multimodal sarcastic tweets via modeling cross-modality contrast and semantic association. Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 3777-3786.

- [30] Yu, W., Xu, H., Yuan, F., Luo, X., Li, J., & Huang, R. (2022). Learning to detect and explain sarcasm for sentiment analysis in social media. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30, 765-775.

Thesis

ORIGINALITY REPORT

5%

SIMILARITY INDEX

3%

INTERNET SOURCES

5%

PUBLICATIONS

3%

STUDENT PAPERS

PRIMARY SOURCES

1

d197for5662m48.cloudfront.net

Internet Source

<1 %

2

www.themoonlight.io

Internet Source

<1 %

3

etds.stu.edu.tw

Internet Source

<1 %

4

Submitted to University of Central Lancashire

Student Paper

<1 %

5

"Pattern Recognition and Computer Vision",
Springer Science and Business Media LLC,
2025

Publication

<1 %

6

Submitted to Dublin Business School

Student Paper

<1 %

7

peerj.com

Internet Source

<1 %

8

Boddu, Sindhu. "Energy-Efficient Object
Detection for Aerial Imagery: Quantization
and Deployment of YOLOv4-Tiny on
Raspberry Pi", The University of North
Carolina at Charlotte

Publication

<1 %

9

Muhammad Yaseen Khan, Tafseer Ahmed,
Muhammad Shoaib Siddiqui, Shaukat Wasi.
"Cognitive Relationship-based Approach for
Urdu Sarcasm and Sentiment Classification",
IEEE Access, 2023

Publication

<1 %

10

Jing Dong, Yu Sui, Qiang Zhang, Hui Fang,
Gerald Schaefer, Rui Liu, Pengfei Yi, Xiaoyong
Fang. "An adaptive multimodal semantic
knowledge enhanced framework for sarcasm
detection", Expert Systems with Applications,
2026

Publication

<1 %

11

Submitted to Purdue University

Student Paper

<1 %

12

proceedings.mlr.press

Internet Source

<1 %

13

"Advances in Soft Computing", Springer
Science and Business Media LLC, 2026

Publication

<1 %

14

"Generative AI in Creative Industries",
Springer Science and Business Media LLC,
2025

Publication

<1 %

15

Mihir Narayan Mohanty, Bibhuprasad
Mohanty, Kandarpa Kumar Sarma, Dmitrii
Kaplun. "Futuristic Information and
Communication: A Multimodal
Multidisciplinary Signal Analysis", CRC Press,
2026

Publication

<1 %

16	Submitted to Queen Mary and Westfield College Student Paper	<1 %
17	B. G. Nagaraja, S. Kannadhasan. "Information and Communication Systems", CRC Press, 2026 Publication	<1 %
18	dtu.ac.in Internet Source	<1 %
19	Submitted to University of Hertfordshire Student Paper	<1 %
20	ejurnal.methodist.ac.id Internet Source	<1 %
21	jutif.if.unsoed.ac.id Internet Source	<1 %
22	A. Punitha, S. Syedakbar, S. Jeyasudha. "Intelligent and Sustainable Systems: AI, Green IoT, and Adaptive Automation in Electrical and Communication Technologies - Proceedings of International Conference in Pathway of Electrical, Automation, Communication, and Electronics (PEACE-2025)", CRC Press, 2026 Publication	<1 %
23	Gyan Prakash, Amandeep Kaur. "AI and Sustainable Transformations", CRC Press, 2025 Publication	<1 %
24	ejournal.uin-malang.ac.id Internet Source	<1 %

25

www.idl.iscram.org

Internet Source

<1 %

26

www.mdpi.com

Internet Source

<1 %

27

repository.effatuniversity.edu.sa

Internet Source

<1 %

Exclude quotes On

Exclude matches < 12 words

Exclude bibliography On