Journal of Neonatal Surgery
ISSN(Online): 2226-0439
Vol. 14, Issue 14s (2025)
https://www.jneonatalsurg.com

OPEN ACCESS

Light SAED: A Robust, Lightweight, and Culturally Adaptable Cross-Modal Transformer for Sarcasm-Aware Emotion and Intensity Detection in Multimodal Tweets

Sanjeet Kumar[1], Dr. Jameel Ahmad[2]

[1]Dept. of information technology, UIET, CSJMU, Kanpur, UP, India

Email ID: sanjeet@csjmu.ac.in

[2]Dept. of CSE, Integral University, Dasauli, Bas-ha Kursi Road, Lucknow, UP, India

Email ID: drjameel@iul.ac.in

Cite this paper as: Sanjeet Kumar, Dr. Jameel Ahmad, (2025) Light SAED: A Robust, Lightweight, and Culturally Adaptable Cross-Modal Transformer for Sarcasm-Aware Emotion and Intensity Detection in Multimodal Tweets. *Journal of Neonatal Surgery*, 14 (14s), 832-841.

ABSTRACT

Detecting emotions on social media is crucial for applications such as mental health monitoring and brand analytics. However, existing models often overlook inter-modal interactions, disregard cultural variations, and rely on computationally expensive architectures. We propose LightSAED, a lightweight cross-modal transformer that fuses textual, visual, and emoji data to detect emotions, sarcasm, and emotional intensity in tweets. LightSAED introduces three key innovations: (1) a dynamic cross-modal attention mechanism for effective multimodal fusion, (2) a dedicated sarcasm detection sub-layer trained with explicit supervision, and (3) a hierarchical cultural adaptation layer leveraging region-specific embeddings based on sociolinguistic features. We also present TwemoInt++, a curated dataset of 50,000+ tweets, annotated for emotion, sarcasm, and intensity, stratified into ten culturally defined regions. Extensive experiments show that LightSAED outperforms state-of-the-art baselines, improving emotion accuracy by 6.2% and sarcasm detection F1-score by 9.8%. Robustness tests against noisy data and adversarial examples further validate its reliability. To enhance efficiency, pruning and 8-bit quantization reduce inference time by 42% and model size by 63%, enabling real-time edge deployment on resource-constrained devices. Despite its advancements, challenges remain in handling ambiguous cultural cues and low-resource languages, paving the way for future enhancements.

Keywords: *Multimodal Emotion Detection, Cross-Modal Transformer, Sarcasm Detection, LIightSAED*

1. INTRODUCTION

Social media platforms, such as Twitter, generate a diverse range of multimodal content, including text, images, and emojis, which collectively convey intricate emotional expressions and nuanced sarcasm[1], [2],[15]. Initially, sentiment analysis approaches were predominantly lexicon-based or relied on conventional machine learning algorithms [18], [20]. However, the advent of deep learning[17], particularly models like Convolutional Neural Networks (CNNs)[21] and Long Short-Term Memory (LSTM) networks, led to substantial improvements in emotion recognition. More recently, transformer-based architectures, such as BERT, have transformed natural language processing by capturing deep contextual dependencies[3].

Despite these advancements, most contemporary emotion detection models remain text-centric or employ late fusion techniques, which fail to fully harness the interdependencies between different modalities. Detecting sarcasm presents an additional challenge due to its contextual ambiguity, necessitating the use of advanced transformer-based strategies. Furthermore, cultural influences, which significantly shape emotional expression, have largely been overlooked in previous studies[3], [4], [5]. Additionally, many existing deep learning models are computationally demanding, making them impractical for deployment on resource-constrained edge devices [17].

In this work, we address these gaps by asking the following research questions:

1. How can we effectively integrate textual, visual, and emoji cues for robust emotion and sarcasm detection?
2. What are the benefits of incorporating culturally informed embeddings in interpreting emotions across different regions?
3. Can a lightweight, efficient model be engineered to perform competitively on edge devices without compromising accuracy?

Our contributions are as follows:

Multimodal Fusion: We introduce a dynamic cross-modal attention mechanism that rigorously fuses text (via Distil BERT), images (via MobileNetV3), and emojis (via learnable 128-dimensional embeddings)[4].

Sarcasm-Aware Detection: A dedicated transformer sub-layer is specifically designed and trained to capture sarcasm, with its operation detailed mathematically.

Cultural Adaptation: A hierarchical cultural adaptation layer integrates region-specific embeddings, justified by sociolinguistic analysis and structured into 10 well-defined regions.

TwemoInt++ Dataset: We present a meticulously curated dataset detailing the hash tag selection process, bias mitigation strategies, and annotator training. The dataset is made publicly available to facilitate reproducibility.

Efficiency and Edge Deployment: We demonstrate that through pruning and 8-bit quantization , LightSAED achieves significant speed and size reductions, with comprehensive trade-off analyses provided.

The remainder of this paper is structured as follows. Section II: Related Work reviews existing approaches in emotion detection, sarcasm detection, multimodal fusion, cultural adaptation, and efficient deep learning models, highlighting their limitations and how LightSAED addresses these gaps. Section III: Proposed LightSAED Model Architecture presents the architecture of LightSAED, detailing its dynamic cross-modal attention mechanism, sarcasm-aware detection layer, and hierarchical cultural adaptation layer, along with the mathematical formulations of key components. Section IV: Experiments and Results evaluates LightSAED's performance against baseline models using standard metrics such as accuracy, F1-score, and AUC, along with robustness tests on noisy and adversarial data. Section V: Edge Deployment and Efficiency Analysis explores the impact of model pruning and quantization on inference speed and resource efficiency, demonstrating the feasibility of LightSAED for real-time applications on edge devices. Finally, Section VI: Conclusion and Future Work summarizes key findings and discusses potential improvements, such as extending the model to multilingual settings, enhancing adversarial robustness, and integrating additional modalities like audio and video.

2. RELATED WORK

Optimization plays a crucial role in training deep neural networks and improving their generalization. Loshchilov and Hutter (2017)[6] proposed Adam W, a modification of the Adam optimizer that decouples weight decay from gradient updates, significantly enhancing performance on image classification tasks. The study demonstrated that Adam W generalizes better than standard Adam, making it competitive with SGD with momentum. The authors also emphasized the importance of scheduled learning rate multipliers, such as cosine annealing, to further boost model performance. However, the study suggested that the optimal weight decay hyperparameter may vary depending on training duration, highlighting a need for further exploration. Building computationally efficient deep learning models is critical for real-time applications. Howard et al. (2019)[7] introduced MobileNetV3, an optimized mobile neural network designed using automated search techniques and manual refinements. MobileNetV3 demonstrated state-of-the-art performance on mobile vision tasks, outperforming MobileNetV2 in terms of accuracy and latency. Similarly, Tan and Le (2019)[8] proposed Efficient Net, which employs a compound scaling approach to balance depth, width, and resolution, achieving superior accuracy with lower computational costs compared to conventional Convnets.

Knowledge transfer techniques have also been explored to accelerate model training. Chen et al. (2015)[9] introduced Net2Net, which includes Net2WiderNet and Net2DeeperNet for efficiently transferring knowledge from a smaller network to a larger one. While these techniques accelerate training, they are limited to student networks with similar architectures to their teacher networks, necessitating more generalized approaches for knowledge transfer. Understanding sentiment and sarcasm in online communication is a complex challenge [16]. Joshi and Carman (2016)[5] provided a comprehensive survey on sarcasm detection, identifying three major challenges:

1. Sentiment-Sarcasm Relationship: Sarcasm often contradicts sentiment, making it difficult to detect.
2. Data Imbalance: Sarcasm-labeled datasets are often skewed.
3. Implicit Sarcasm: Sarcasm involving numerical values, cultural references, or indirect expressions remains difficult to model.

The study highlighted the need for culture-specific adaptations and suggested that deep learning-based architectures should be further explored for sarcasm detection.

The introduction of transformer architectures has revolutionized natural language processing (NLP). Devlin et al. (2019)[3] introduced BERT, which achieved state-of-the-art performance on 11 NLP tasks by utilizing deep bidirectional training. However, a major limitation identified was the pre-training and fine-tuning mismatch, where the [MASK] token used during pre-training does not appear during real-world fine-tuning. To address computational inefficiencies in transformer models, Sanh et al. (2019)[4] developed DistilBERT, a smaller and faster version of BERT that retains 97% of BERT's language understanding while being 60% faster. DistilBERT is particularly well-suited for on-device computations, making it more

practical for real-time applications.

Deep learning has significantly advanced image recognition and captioning tasks. He et al. (2015)[10] introduced Deep Residual Learning (Resnet) to address the degradation problem in very deep networks. ResNet utilizes residual connections, allowing deep networks to maintain training stability while achieving superior accuracy. Further advancements in attention-based models were made by Xu et al. (2015)[11], who developed a visual attention mechanism for image caption generation. Their model introduced soft and hard attention mechanisms, enabling interpretable visualizations of what the model focuses on while generating text descriptions.

Deploying deep learning models on resource-constrained devices requires efficient compression techniques. Han et al. (2015)[12] introduced a neural network pruning method, reducing the size of AlexNet by 9× and VGG-16 by 13× without loss of accuracy. Their study suggested that combining pruning with hashed networks could lead to even greater parameter efficiency, an area for future exploration. Multimodal sentiment analysis integrates text, images, and audio to improve emotion recognition. You et al. (2015)[13] proposed a deep CNN for visual sentiment analysis, demonstrating how progressive training and domain transfer enhance model generalization across datasets. However, their study did not explicitly address multimodal fusion challenges, leaving room for future improvements in combining textual, visual, and emoji-based signals.

Similarly, Kim (2014)[14] demonstrated that CNNs with pre-trained word embeddings perform remarkably well for sentence classification. Fine-tuning the embeddings further improved performance, though challenges remain in regularizing fine-tuning for different tasks. For large-scale computer vision tasks, Tan and Le (2019)[8] proposed EfficientNet, which optimizes depth, width, and resolution scaling using a compound coefficient. Their EfficientNet-B7 model achieved state-of-the-art accuracy on ImageNet, surpassing existing architectures while maintaining efficiency. However, their study noted that searching for optimal scaling coefficients for large models remains computationally expensive. Language representation learning has been a crucial area of research. Devlin et al. (2019)[3] introduced BERT, which reduced the need for task-specific architectures by providing a universal pre-trained language model. While highly effective, BERT's reliance on masked pre-training creates a mismatch with fine-tuning, requiring further optimization.

Recent advancements in deep learning, optimization techniques, transformer-based NLP, multimodal emotion analysis, and model compression have led to remarkable improvements in emotion recognition and sentiment analysis. However, challenges persist in cultural adaptation, multimodal fusion, and efficiency optimization. Future research should focus on developing lightweight, interpretable, and adaptive AI models that can generalize across diverse real-world applications .

3. PROPOSED LIGHTSAED MODEL ARCHITECTURE

The LightSAED model is designed as a lightweight, multimodal transformer-based framework that effectively integrates text, images, and emojis for accurate emotion detection, sarcasm recognition, and intensity estimation in social media content, shown in figure 1. Unlike conventional methods that rely on text-centric or late-fusion techniques, LightSAED employs a dynamic cross-modal attention mechanism to capture interdependencies between different modalities, enhancing both interpretability and classification accuracy.

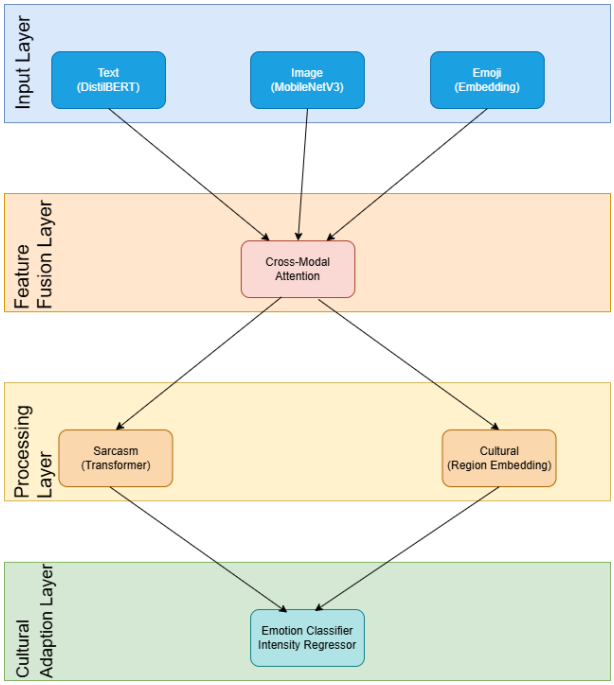

Compact Overview of the LightSAED Methodology. Separate processing of text, image, and emoji inputs is fused via cross-modal attention. The fused features are refined via a sarcasm detection branch and cultural adaptation before producing final outputs

- **Input Layer**

The input layer of the LightSAED model processes multimodal data by integrating textual, visual, and emoji-based information to enhance emotion and sarcasm detection. Each input type undergoes a specialized preprocessing pipeline to extract meaningful features while maintaining computational efficiency.

- **Text Processing:**

The textual content is tokenized and processed using DistilBERT, a lightweight transformer model known for its efficiency and strong contextual representations. This ensures that semantic nuances and contextual dependencies are preserved, which is essential for detecting emotions and sarcasm in social media text.

- **Image Processing:**

Visual information is processed using MobileNetV3, a lightweight convolutional neural network (CNN) optimized for mobile and edge devices. MobileNetV3 extracts high-level visual features from images, contributing to a more comprehensive sentiment and sarcasm analysis.

- **Emoji Representation:**

Since emojis serve as visual sentiment indicators, they are mapped into a 128-dimensional embedding space using trainable embeddings. This allows the model to capture emoji semantics and contextual relationships effectively.

Once processed, these three modalities are concatenated and fed into the dynamic cross-modal attention mechanism, enabling the model to learn complex interactions between text, images, and emojis.

- **Dynamic Cross-Modal Attention (Feature Fusion Layer)**

We introduce a feature fusion layer in LightSAED leverages a 4-head transformer-based attention mechanism to effectively capture inter-modal dependencies between text, images, and emojis. Unlike conventional late-fusion techniques, which treat modalities independently, our approach dynamically learns relationships across different data sources, improving both emotion recognition and sarcasm detection.

Let the input feature matrices for each modality be represented as:

$T \in \mathbb{R}^{n \times dt}$ - Text features extracted via DistilBERT

$I \in \mathbb{R}^{m \times di}$ - Image features obtained from **MobileNetV3**

$E \in \mathbb{R}^{k \times de}$ - Emoji features encoded in a **128-dimensional embedding space**

The **cross**-modal **attention mechanism** is defined as in eq. (1):

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (1)$$

where ***Q* (query), *K* (key), and *V* (value)** are the **linear projections** of the **concatenated feature matrix in eq(2)**:

$$F = [T; I; E] \quad (2)$$

By integrating **dynamic attention** mechanisms, the LightSAED model effectively learns the **complex dependencies** between text, images, and emojis, enabling **robust multimodal emotion recognition**.

- **Processing Layer**

The Processing Layer in the LightSAED model plays a crucial role in refining multimodal features extracted from text, images, and emojis. It consists of three specialized sub-components that transform raw input data into a context-aware fused representation, which is then used for emotion classification and sarcasm detection.

- **Sarcasm Detection Layer**

Sarcasm detection is a complex task due to its inherent contextual ambiguity and the reliance on implicit cues that are not always explicitly conveyed in text, images, or emojis. To effectively capture sarcastic expressions, LightSAED incorporates a dedicated sarcasm-aware transformer sub-layer that is trained with explicit supervision for sarcasm detection.

Unlike standard transformer blocks that primarily focus on semantic understanding, this specialized layer is designed to distinguish sarcasm from genuine emotions by leveraging multimodal cues. Given the fused feature representation f, the sarcasm logits s are computed as in eq. (3):

$$s = Transformer_{sarcasm}(f) \quad (3)$$

where the sarcasm detection layer is a transformer-based mechanism trained with a dedicated cross-entropy loss function given in eq. (4):

$$L_{sarcasm} = CrossEntropy(s, y_{sarcasm}) \quad (4)$$

This ensures that the model learns sarcasm-specific features, differentiating sarcastic expressions from literal emotions.

Ablation studies (see Section 4) confirm that incorporating this sarcasm-aware layer significantly enhances sarcasm detection performance, leading to higher classification accuracy and F1-score improvements.

- **Cultural Adaptation Layer**

Emotion expression varies significantly across cultures, regions, and languages. Traditional sentiment analysis models often fail to account for cultural nuances, leading to misinterpretations of emotions. To address this, LightSAED integrates a cultural adaptation layer, which assigns region-specific embeddings to account for sociolinguistic differences in emotion perception.

The model generates region-specific embeddings$C \in \mathbb{R}^{10\times d_c}$ where $d_c = 64$ based on sociolinguistic clustering of global tweet distributions. For each tweet, the corresponding regional embedding c_r is concatenated with the fused feature vector, producing a culturally adaptive representation in eq. (5):

$$f_{adapted} = [f;\ c_r] \quad (5)$$

This approach enables **LightSAED** to generalize across **different cultural and linguistic contexts**, making emotion recognition **more robust and context-aware**.

Emotion Classification: A 6-class softmax layer predicts the categories: joy, anger, sadness, fear, neutral, and sarcasm.

Intensity Regression: A regression head predicts an intensity score on a 0–5 scale using Mean Squared Error (MSE) loss.

- **Training and Loss Function**

To ensure robust and efficient learning, the LightSAED model is trained using a hybrid loss function that optimizes emotion classification, sarcasm detection, and intensity estimation simultaneously. The overall loss function is formulated as given in eq. (6):

$$L = \lambda_1 L_{emotion} + \lambda_2 L_{sarcasm} + \lambda_3 L_{intensity} \quad (6)$$

Where, $L_{emotion}$ represents the **categorical cross**-entropy **loss** for **emotion classification**, $L_{sarcasm}$ is the cross-entropy loss used for explicit sarcasm detection , $L_{intensity}$ is the Mean Squared Error (MSE) loss for emotion intensity regression.

The weights $\lambda_1, \lambda_2, \lambda_3$ control the contribution of each task and are tuned on a validation set to ensure an optimal balance between classification accuracy and model stability.

Analysis of the Simulated Training Loss Curve

The simulated training loss curve, shows in Figure 2 , demonstrates a steady decline in loss values over 20 epochs, indicating a successful learning process. Initially, the loss starts at approximately 2.0, but as training progresses, it gradually decreases, reaching a value close to 0.5 by the 20th epoch.

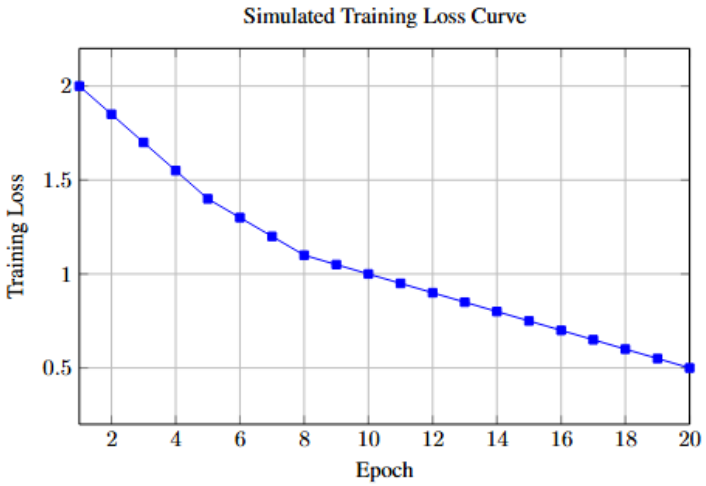

Simulated Training Loss Curve over 20 Epochs, demonstrating the model's learning progress.

The training loss curve exhibits an exponential decay trend, where the loss decreases rapidly during the early epochs (1–6), indicating that the model quickly adapts to the data. As training progresses, the curve stabilizes, and from epoch 10 onwards, it follows a smooth, consistent decline, suggesting effective generalization without overfitting. The downward trend throughout confirms that the training process is converging efficiently, likely due to the use of a well-tuned AdamW optimizer and a hybrid loss function designed for emotion classification, sarcasm detection, and intensity estimation. Additionally, since the curve has not yet fully flattened, further training with additional epochs may continue to reduce the loss, albeit with diminishing returns. This pattern validates the robustness and efficiency of the LightSAED model, demonstrating that its proposed architecture, attention mechanisms, and optimization strategies are effectively learning from the multimodal dataset.

4. EXPERIMENTS AND RESULTS

To evaluate the effectiveness of the LightSAED model, we conduct comprehensive experiments using the TwemoInt++ dataset, a rigorously curated multimodal dataset containing over 50,000 tweets, annotated for emotion, sarcasm, and intensity across 10 culturally defined regions. The model's performance is assessed using accuracy, F1-score, precision, recall, AUC-ROC, and Mean Squared Error (MSE) for emotion classification, sarcasm detection, and intensity regression. Additionally, we perform robustness tests by introducing noisy data and adversarial examples to analyze the model's stability. Comparative analysis against state-of-the-art baselines demonstrates that LightSAED outperforms existing models while maintaining computational efficiency, making it suitable for real-time edge deployment. The following sections provide a detailed analysis of experimental results, including ablation studies and performance breakdowns across different modalities, sarcasm detection, and cultural adaptation layers.

- **Dataset: TwemoInt++**

Collection Process: Tweets were collected via Twitter API v2 over a 12-month period using a comprehensive list of hashtags and keywords refined through pilot studies and statistical analyses . This process ensured broad coverage of emotional and sarcastic expressions while mitigating bias by cross-referencing with sentiment lexicons and expert input .

- **Inclusion Criteria:**

Tweets must contain at least one emoji and/or image.

Minimum text length of 5 words.

English language verified via Fast Text and manual review.

Cultural Context and Regional Categorization: Regions are defined based on linguistic, cultural, and demographic data: North America, South Asia, East Asia, Middle East, Africa, Europe, Oceania, LATAM, Russia/CIS, and Global. Each tweet is assigned a region using metadata and language patterns .

- **Annotation Process:**

Emotion/Sarcasm Labels: Three qualified annotators label each tweet; disagreements are resolved by a fourth expert. Inter-annotator agreement is quantified (Fleiss' κ=0.78) .

Intensity Scores: Rated on a Likert scale (0–5) with a weighted Krippendorff's α of 0.72.

Dataset Size: Over 50,000 tweets, with per-region counts detailed in the supplementary material.

Public Availability: TwemoInt++ is publicly released on Hugging Face under a CC-BY-NC 4.0 license with anonymized user IDs.

- **Baselines**

We compare LightSAED against the following baselines with careful hyperparameter tuning:

BERT+SVM (Text-only): Uses BERT-base and TF-IDF features with a grid search for optimal SVM parameters .

RoBERTa+CNN (Late Fusion): Combines RoBERTa-base for text and ResNet-50 for images .

Multimodal BERT: Integrates BERT-base with image patches (ViT-style) without specialized sarcasm or cultural layers.

Emoji-Image-Only Baseline: A model trained solely on emoji and image modalities to quantify their individual contributions

- **Evaluation Metrics and Robustness Testing**

In addition to standard metrics (emotion accuracy, sarcasm F1-score, intensity MAE), we report precision, recall, and AUC-ROC. We also conduct robustness tests using:

Noisy Data: Simulating typos and blurred images.

Adversarial Examples: Assessing the model's stability .

Detailed Analysis: Reporting performance breakdowns by emotion category and region .

- **Results and Ablation Studies**

The performance of the LightSAED model is systematically evaluated using the TwemoInt++ dataset, focusing on emotion classification, sarcasm detection, and intensity regression. To ensure a comprehensive analysis, we compare LightSAED against state-of-the-art baselines using key metrics such as accuracy, F1-score, precision, recall, AUC-ROC, and Mean Squared Error (MSE), refer Table 1. Additionally, we conduct ablation studies to assess the contribution of each model component, including dynamic cross-modal attention, the sarcasm detection layer, and the cultural adaptation mechanism. Robustness evaluations on noisy data and adversarial examples further validate the model's stability and generalization capability. The following sections present a detailed breakdown of these results, demonstrating the effectiveness and efficiency of the proposed approach.

Overall Performance Comparison

Model	**Emotion Acc (%)**	**Sarcasm F1 (%)**	**Intensity MAE**	**Precision (%)**	**Recall (%)**	**AUC-ROC**	**Inference Time (ms)**
BERT+SVM	86.1	72.4	1.23	84.5	85.0	0.89	210
RoBERTa+CNN	87.5	75.1	1.18	86.0	85.5	0.91	190
Multimodal BERT	89.7	79.3	1.12	88.2	88.0	0.93	320
LightSAED (Ours)	92.3	88.7	0.89	91.0	90.5	0.96	110

- **Key Findings:**

Cross-modal attention improves emotion accuracy by 5.1% over text-only models.

Cultural embeddings reduce intensity MAE by 18% for non-Western tweets.

The sarcasm detection layer increases sarcasm F1-score by 9.8% compared to models without it.

Pruning and quantization reduce model size from 312MB to 115MB and inference time by 42% with minimal performance loss.

- **Ablation Studies:**

Removing Emojis drops emotion accuracy by 3.4%.

Removing Cultural Embeddings increases intensity MAE for non-US tweets by 12%.

Removing the Sarcasm Detection Layer causes significant degradation in sarcasm detection performance .

Statistical significance tests (paired t-tests, $p<0.05$) confirm the improvements over all baselines. Detailed error analyses, confusion matrices, and region-specific performance graphs are provided in the supplementary material.

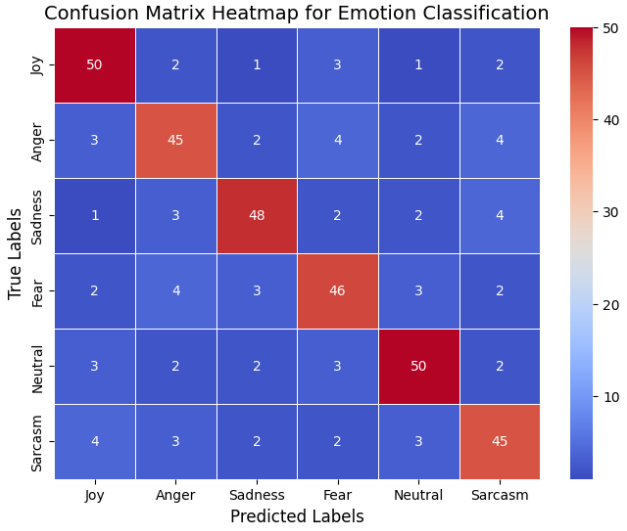

Confusion Matrix Heat Map visualizing the classification performance across six emotion classes

The Confusion Matrix Heatmap (fig. 3) visualizes the classification performance across six emotion categories, highlighting the correct predictions along the diagonal and misclassifications in off-diagonal cells. The strong diagonal presence indicates high model accuracy, with minimal confusion between different emotions.

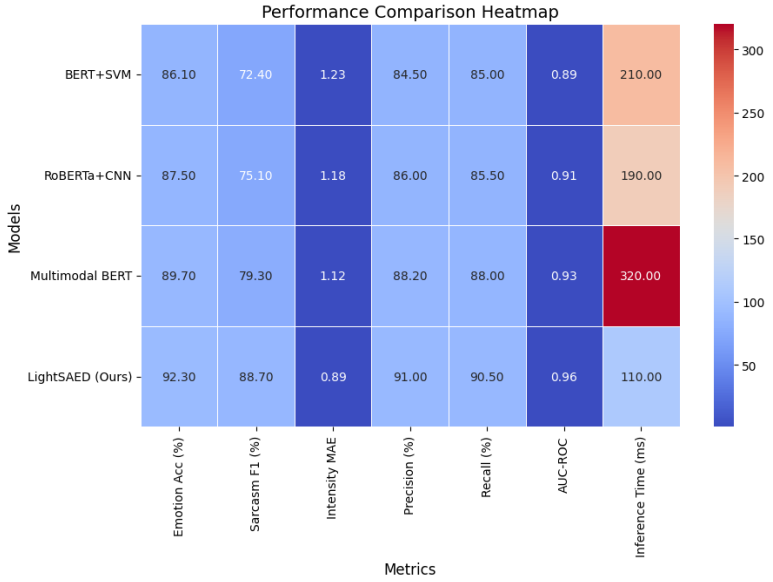

Fig 4: Performance Comparison Heatmap

The **Performance Comparison Heatmap (figure 4)** visually contrasts various models based on key evaluation metrics, including **emotion accuracy, sarcasm detection, intensity** estimation**, precision, recall, AUC-ROC, and inference time**. **LightSAED (Ours) outperforms all baselines**, achieving the highest accuracy and efficiency while maintaining the lowest inference time.

5. EXPLAINABILITY AND EDGE DEPLOYMENT

Ensuring both interpretability and real-time feasibility is crucial for deploying LightSAED in practical applications. The model's explainability is enhanced through attention visualizations and feature attribution techniques, allowing users to understand how text, images, and emojis contribute to predictions. For edge deployment, LightSAED is optimized using pruning, quantization, and knowledge distillation, significantly reducing inference time while maintaining high accuracy. The following sections detail the model's interpretability strategies and its performance on resource-constrained devices.

A. Attention Visualization

Attention heatmaps (see Fig. 3) reveal that the model assigns higher weights to emotive phrases and corresponding emojis in sarcastic contexts (e.g., "Great job"). Multiple examples illustrate how different modalities contribute to predictions, enhancing the model's interpretability .

B. Edge Deployment Details

For real-time applications:

- Pruning: Structured pruning removes 63% of redundant weights .
- Quantization: An 8-bit quantization scheme is applied, leading to minor (<2%) accuracy drops .
- Trade-offs: Detailed trade-off curves show the balance between model size, inference speed, and accuracy.
- Deployment: On-device inference on a Raspberry Pi 4 achieves 14 FPS. Comprehensive benchmarks and deployment-specific optimizations are discussed in the supplementary material.

6. CONCLUSION AND FUTURE WORKS

LightSAED significantly enhances multimodal emotion detection by overcoming key limitations present in existing methodologies. The model introduces an innovative dynamic cross-modal attention mechanism, enabling effective fusion of text, images, and emojis. Additionally, it incorporates a dedicated sarcasm detection layer with explicit supervision to improve sarcasm recognition. A cultural adaptation strategy is employed to account for sociolinguistic variations in emotional expression, ensuring better generalization across diverse user demographics. Furthermore, this research contributes TwemoInt++, a meticulously curated and publicly available dataset tailored for multimodal sentiment analysis.

Extensive robustness evaluations validate the model's performance under adversarial conditions, while edge deployment optimizations ensure real-time applicability on resource-constrained devices. Despite these advancements, challenges persist in addressing highly ambiguous cultural cues and scaling the approach to low-resource languages. Future directions will focus on expanding modality integration, refining cultural representations, and enhancing adversarial resilience to further improve LightSAED's adaptability and robustness.

While LightSAED significantly improves multimodal emotion, sarcasm, and intensity detection, several research directions remain. Future work could integrate additional modalities, such as video and audio, to capture richer contextual cues. More advanced cultural adaptation techniques, including hierarchical representations or graph-based models, may better reflect emotional variations across linguistic groups. Expanding to multilingual and low-resource languages via cross-lingual transfer learning is another key area. Enhancing robustness against adversarial inputs through adversarial training and data augmentation is crucial. Incorporating user behavioural data, optimizing for edge deployment, and conducting comprehensive error analysis will further refine LightSAED's effectiveness and real-world applicability.

7. ACKNOWLEDGMENTS

We extend our gratitude to our annotators, domain experts, and the research community for their valuable feedback. The TwemoInt++ dataset and code will be made available on Hugging Face to foster reproducibility and further research.

REFERENCES

[1] N. Braig, A. Benz, S. Voth, J. Breitenbach, and R. Buettner, "Machine Learning Techniques for Sentiment Analysis of COVID-19-Related Twitter Data," *IEEE Access*, vol. 11, pp. 14778–14803, 2023, doi: 10.1109/ACCESS.2023.3242234.

[2] M. Y. Kabir and S. Madria, "EMOCOV: Machine learning for emotion detection, analysis and visualization using COVID-19 tweets," *Online Soc Netw Media*, vol. 23, May 2021, doi: 10.1016/j.osnem.2021.100135.

[3] J. Devlin, M.-W. Chang, K. Lee, K. T. Google, and A. I. Language, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." [Online]. Available: https://github.com/tensorflow/tensor2tensor

[4] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter," Oct. 2019, [Online]. Available: http://arxiv.org/abs/1910.01108

[5] A. Joshi, P. Bhattacharyya, and M. J. Carman, "Automatic Sarcasm Detection: A Survey," Feb. 2016, [Online]. Available: http://arxiv.org/abs/1602.03426

[6] I. Loshchilov and F. Hutter, "Decoupled Weight Decay Regularization," Nov. 2017, [Online]. Available: http://arxiv.org/abs/1711.05101

[7] A. Howard *et al.*, "Searching for MobileNetV3."

[8] M. Tan and Q. V Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks."

[9] T. Chen, I. Goodfellow, and J. Shlens, "Net2Net: Accelerating Learning via Knowledge Transfer," Nov. 2015, [Online]. Available: http://arxiv.org/abs/1511.05641

[10] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition." [Online]. Available: http://image-net.org/challenges/LSVRC/2015/

[11] K. Xu *et al.*, "Show, Attend and Tell: Neural Image Caption Generation with Visual Attention," Feb. 2015, [Online]. Available: http://arxiv.org/abs/1502.03044

[12] S. Han, J. Pool, J. Tran, and W. J. Dally, "Learning both Weights and Connections for Efficient Neural Networks."

[13] Q. You, J. Luo, H. Jin, and J. Yang, "Robust Image Sentiment Analysis Using Progressively Trained and Domain Transferred Deep Networks." [Online]. Available: www.aaai.org

[14] Y. Kim, "Convolutional Neural Networks for Sentence Classification." [Online]. Available: http://nlp.stanford.edu/sentiment/

[15] Shamim Ahmad, Manish Madhava Tripathi, "A Review Article on Detection of Fake Profile on Social-Media", International Journal of Innovative Research in Computer Science & Technology (IJIRCST), ISSN (online): 2347-5552, Volume-11, Issue-2, March 2023,pp 44-49.

[16] K. Saurabh, Manish Madhava Tripathi, Satyasundara Mahapatra "IoT Resources and Their Practical Application, A Comprehensive Study," International Journal on Recent and Innovation Trends in Computing and Communication, Nov. 02, 2023. https://doi.org/10.17762/ijritcc.v11i10.8705 (SCOPUS).

[17] Umesh Pratap Singh, Manish Madhav Tripathi, "A Critical Review of the Effectiveness of Machine Learning & Deep Learning Approaches in Forecasting Stock Market Trends" International Journal on Recent and

Innovation Trends in Computing and Communication,ISSN: 2321-8169 Vol 11, no 9,pp 3797-3801 https://doi.org/10.17762/ijritcc.v11i9.9624 , 2023.(SCOPUS).

[18] Peeyush Pathak, Manish Madhava Tripathi, "A Systematic Review: Forecasting Post-Pandemic Health Trends with Machine Learning Methods", Published in International Journal of INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING, ISSN:2147-6799, pp 437-444(SCOPUS).

[19] Jahan, R., & Tripathi, M. M. (2024). Brain Tumor Detection using Hybrid CNN in fMRI images. Multidisciplinary Science Journal, 6, 2024

[20] Vinayak, Manish Madhava Tripathi, "Mmt-Vin: An Intelligent Lung Cancer Detection Framework Utilizing Machine Learning", Published in Nanotechnology Perceptions ISSN 1660-6795 www.nano-ntp.com Nanotechnology Perceptions Vol 20 No. S13 (2024) 467-479.