



AVR: Synergizing Foundation Models for Audio-Visual Humor Detection

Sarthak Sharma^{*1}, Orchid Chetia Phukan^{*1}, Drishti Singh¹, Arun Balaji Buduru¹, Rajesh Sharma^{1,2}

¹IIT-Delhi, India, ²University of Tartu, Estonia

^{*}equal contribution

sarthak23083@iiitd.ac.in , orchidp@iiitd.ac.in

Abstract

In this work, we present, **AVR** application for audio-visual humor detection. While humor detection has traditionally centered around textual analysis, recent advancements have spotlighted multimodal approaches. However, these methods lean on textual cues as a modality, necessitating the use of ASR systems for transcribing the audio-data. This heavy reliance on ASR accuracy can pose challenges in real-world applications. To address this bottleneck, we propose an innovative audio-visual humor detection system that circumvents textual reliance, eliminating the need for ASR models. Instead, the proposed approach hinges on the intricate interplay between audio and visual content for effective humor detection.

Index Terms: Audio-Visual Humor Detection, Multimodal System, VideoMAE, Audio Spectrogram Transformer, Languagebind

1. Introduction

Humor has intrigued researchers since ancient Greek times [1], and numerous theories have been proposed since then. Theories have suggested that humor encompasses a combination of semantic context, facial expressions, gestures, and voice modulation, rendering it complex for even humans to detect in various situations [2]. These challenges have spurred researchers to delve into humor detection, making it a crucial domain in human-computer interaction (HCI).

The exploration of humor detection commenced with the analysis of textual content around 2004 [3], primarily focusing on jokes and puns characterized by word-play and categorizing such sentences as humorous or not. Initially, NLP techniques such as N-grams were employed to assess the similarity scores for predicting the humor quotient of sentences. Subsequently, the scope of datasets expanded to include data from diverse sources like social media platforms such as Twitter [4] and Reddit. Machine learning models such as Support Vector Machine (SVM) and deep learning models such as CNNs and LSTMs were utilized for prediction purposes. To further enhance humor detection, modalities like audio and video were introduced to incorporate facial expressions and voice modulations [5]. TV sitcoms such as "The Big Bang Theory" [6] and Friends have been frequently employed as datasets for multimodal humor analysis due to their incorporation of various humor types such as sarcasm, puns, and satire. Another study [7] utilized the novel Passau-Spontaneous Football Coach Humour (Passau-SFCH) dataset, and trained the models on each modality - speech, text and video and got promising results for the multimodal approach.

The reliance on textual cues for multimodal humor detection systems, which often require accurate transcription through

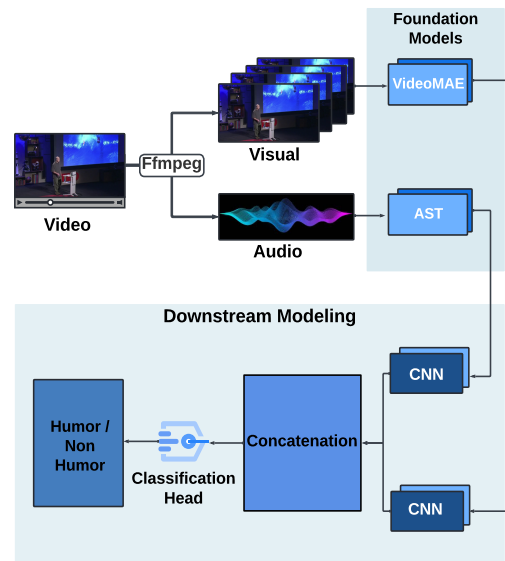


Figure 1: Architecture of the Proposed Modeling Network

ASR systems, can indeed present challenges, particularly in real-world applications where ASR accuracy may vary. To overcome this, we exclusively fuse audio-visual information for more reliable humor detection, removing the need for ASR models. Leveraging state-of-the-art foundation models like VideoMAE, AST, and Languagebind, we extract meaningful audio-visual features for effective humor detection and we name our application as **AVR**. **AVR** intakes a whole video and provides the output without the need for analyzing individual frames in a video.

2. AVR

In this section, we delve into various parts of **AVR** application. First, the foundation models, followed by the downstream modeling, experimental analysis, and lastly the user interface.

Foundation Models: We use Languagebind [8], which is a multimodal model trained for multiple modalities with language as the binding modality. Languagebind shows state-of-the-art (SOTA) performance in various tasks across different modalities. We also use VideoMAE [9], AST [10] which are foundation models trained for video, audio representation learning in self-supervised and supervised fashion. We extract audio-visual representations of 768-dimension from the foundation models. The audios are resampled to 16KHz before passing it to the

Table 1: Evaluation Scores; Scores are average of 5-folds; F1 stands macro-average f1-score

Foundation Models	Accuracy
CNN	
LanguageBind	49.68
VideoMAE + AST	56.70
LSTM	
LanguageBind	48.40
VideoMAE + AST	54.82

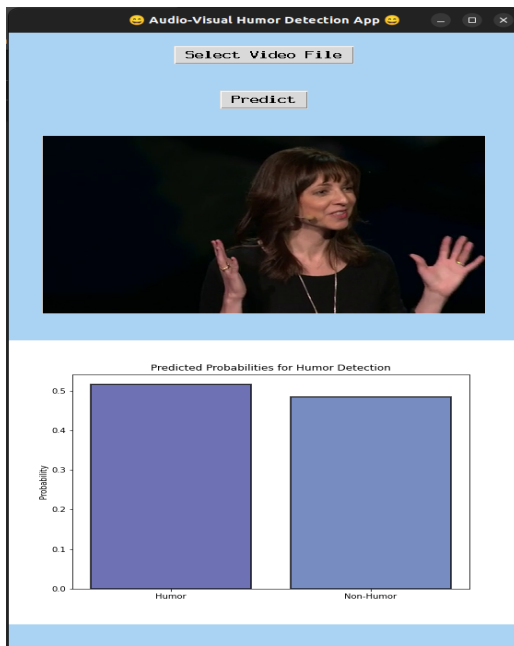


Figure 2: User Interface

foundation models.

Downstream Modelling: We have employed CNN and LSTM as downstream networks. The modeling architecture with CNN as downstream is shown in Figure 1. The CNN model consists of two 1D-CNN with 32 and 64 filters and 3 as filter size. Followed by concatenation for fusion and lastly, classification head with a hidden layer of 128 neurons followed by the output layer. For LSTM model, we use a single LSTM layer with 50 as intermediate followed by the same classification head and the corresponding layer. The models are trained for 50 epochs with a learning rate of $1e-5$. We also employed dropout and early stopping for reducing overfitting. We train the models in a 5-fold manner with 4 folds as training set and one fold as test set. We use cross-entropy as loss function and Adam as the optimizer. The results of the models are given in Table 1. We can see that CNN with VideoMAE and AST representations earned the top-most score. One notable thing is the performance of VideoMAE and AST representations compared to LanguageBind despite it being jointly trained on multiple modalities.

2.1. User Interface

To facilitate efficient user engagement with the model, we have developed an interactive User Interface (UI) leveraging *Tkinter* library. **AVR** provides the benefit of cross-system usage across macOS, Linux, windows, and so on. The UI of **AVR** is shown in Figure 2. Users are prompted to select a .mp4 video file from their local device to upload it. On uploading, the video is segmented into audio and visual modality using *FFmpeg*. Similar to the model training, representations are generated for both the audio and video using AST and VideoMAE respectively and passed to CNN (we use CNN in backend of **AVR** as it shows the best performance). Users are presented with the predictions with a graphical representation in the form of a bar graph, illustrating the probability of the input video being categorized as humor or non-humor. **AVR** takes approx 5-10 secs for inference for videos of duration of 2-30 secs in a P100 GPU.

3. Conclusion

We present **AVR**, an innovative application for audio-visual humor detection using foundation models. This synergy of audio and visual modalities for effective humor detection and thus prevents the usage of text as a modality in multimodal humor detection and in return reducing the dependency on ASR system accuracy.

4. References

- [1] C. Larkin-Galiñanes, “An overview of humor theory,” *The Routledge handbook of language and humor*, pp. 4–16, 2017.
- [2] K. Weber, H. Ritschel, I. Aslan, F. Lingenfeller, and E. André, “How to shape the humor of a robot-social behavior adaptation based on reinforcement learning,” in *Proceedings of the 20th ACM international conference on multimodal interaction*, 2018, pp. 154–162.
- [3] J. M. Taylor and L. J. Mazlack, “Computationally recognizing wordplay in jokes,” in *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 26, no. 26, 2004.
- [4] P. Potash, A. Romanov, and A. Rumshisky, “Semeval-2017 task 6:# hashtagwars: Learning a sense of humor,” in *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, 2017, pp. 49–57.
- [5] O. Weller and K. Seppi, “Humor detection: A transformer gets the last laugh,” *arXiv preprint arXiv:1909.00252*, 2019.
- [6] D. Bertero and P. Fung, “Deep learning of audio and language features for humor prediction,” in *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, 2016, pp. 496–501.
- [7] L. Christ, S. Amiriparian, A. Kathan, N. Müller, A. König, and B. W. Schuller, “Towards multimodal prediction of spontaneous humour: A novel dataset and first results,” *arXiv preprint arXiv:2209.14272*, 2022.
- [8] B. Zhu, B. Lin, M. Ning, Y. Yan, J. Cui, H. Wang, Y. Pang, W. Jiang, J. Zhang, Z. Li *et al.*, “Languagebind: Extending video-language pretraining to n-modality by language-based semantic alignment,” *arXiv preprint arXiv:2310.01852*, 2023.
- [9] Z. Tong, Y. Song, J. Wang, and L. Wang, “Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training,” *Advances in neural information processing systems*, vol. 35, pp. 10 078–10 093, 2022.
- [10] Y. Gong, Y.-A. Chung, and J. Glass, “AST: Audio Spectrogram Transformer,” in *Proc. Interspeech 2021*, 2021, pp. 571–575.