

MM-HSD: Multi-Modal Hate Speech Detection in Videos

Berta Céspedes-Sarrias*

EPFL

Lausanne, Switzerland

Idiap Research Institute

Martigny, Switzerland

berta.cespedessarrias@epfl.ch

Carlos Collado-Capell*

EPFL

Lausanne, Switzerland

Idiap Research Institute

Martigny, Switzerland

carlos.colladocapell@epfl.ch

Pablo Rodenas-Ruiz*

EPFL

Lausanne, Switzerland

Idiap Research Institute

Martigny, Switzerland

pablo.rodenasruiz@epfl.ch

Olena Hrynenko

EPFL

Lausanne, Switzerland

Idiap Research Institute

Martigny, Switzerland

olena.hrynenko@epfl.ch

Andrea Cavallaro

EPFL

Lausanne, Switzerland

Idiap Research Institute

Martigny, Switzerland

andrea.cavallaro@epfl.ch

ABSTRACT

While hate speech detection (HSD) has been extensively studied in text, existing multi-modal approaches remain limited, particularly in videos. As modalities are not always individually informative, simple fusion methods fail to fully capture inter-modal dependencies. Moreover, previous work often omits relevant modalities such as on-screen text and audio, which may contain subtle hateful content and thus provide essential cues, both individually and in combination with others. In this paper, we present MM-HSD, a multi-modal model for HSD in videos that integrates video frames, audio, and text derived from speech transcripts and from frames (i.e. on-screen text) together with features extracted by Cross-Modal Attention (CMA). We are the first to use CMA as an early feature extractor for HSD in videos, to systematically compare query/key configurations, and to evaluate the interactions between different modalities in the CMA block. Our approach leads to improved performance when on-screen text is used as a query and the rest of the modalities serve as a key. Experiments on the HateMM dataset show that MM-HSD outperforms state-of-the-art methods on M-F1 score (0.874), using concatenation of transcript, audio, video, on-screen text, and CMA for feature extraction on raw embeddings of the modalities. The code is available at <https://github.com/idiap/mm-hsd>.

Warning: some of the elements of the paper contain hate speech examples, which could be disturbing to some readers.

KEYWORDS

Hate Speech, Multi-modal fusion, Attention, Social Media.

1 INTRODUCTION

Hate speech (HS) is "*a speech or address inciting hatred or intolerance, especially towards a particular social group on the basis of ethnicity, religious beliefs, sexuality, etc.*" [39]. The widespread use of social media and online fora [35], where people express their opinions on diverse subjects, has led to an increase in HS online [69]. This proliferation of hate-related posts distorts political discourse,

negatively affects public dialogue [50], and can lead to the radicalization of individuals, increasing the risk of hate-related terrorism [45]. Historically, hate speech detection (HSD) was performed manually [45], limiting its scalability and imposing a significant psychological burden on the moderators [68]. With online content growing in volume and complexity, automated HSD reduces the need for human moderation by serving as an initial filter [23, 59].

Since online HS has traditionally been associated with textual content [30], text-based HSD has been extensively studied [8, 28, 48]. However, social media content is increasingly multi-modal, and HS can appear not only in text but also in visuals and audio. To account for this diversity, recent work in multi-modal HSD incorporates multiple data sources, such as images [62, 72, 73] and user metadata [9, 45]. While interest in multi-modal HSD is growing, it remains relatively underexplored compared to text-based approaches. Furthermore, most research on multi-modal HSD has focused on integrating images as an additional modality – particularly in the context of memes and social media posts [8, 22, 31, 53, 63]. In contrast, work on multi-modal detection of HS in videos is relatively scarce, despite the rise of video-centric platforms like YouTube, Instagram, and TikTok, which facilitate its spread.

Video-based HSD is particularly challenging because hateful content may be embedded in multiple modalities, including video frames, on-screen text, and audio [10, 12], often concealed within memes, music, or other non-traditional formats [33]. Some prior work relies solely on audio transcripts [1, 69], overlooking other modalities. Other studies integrate audio, video frames, and transcripts [12, 40, 67], but ignore the visual text within the frames, which could provide significant cues to improve detection accuracy. To our knowledge, there is only one study that incorporates on-screen text as a modality in HSD [71].

Motivated by its strong performance in video classification [24, 55], in this paper, we explore Cross-Modal Attention (CMA) [11, 46] as a flexible and context-aware fusion mechanism. Our insight is that, as CMA allows one modality to attend to another, it is especially useful for identifying necessary contextual cues when HS appears in a different modality. To this end, we propose to use CMA as a feature extractor for multi-modal integration at an early stage.

*These authors contributed equally to this research.

Our main contributions are:

- **Multi-modal HSD in videos:** We contribute to the limited literature on video HSD, developing a model that processes and fuses multiple features across modalities – transcript, audio, video frames, and on-screen text. This multi-modal approach enables a more comprehensive representation of hateful content, especially when individual modalities alone are insufficient or ambiguous.
- **CMA as early fusion for contextual integration:** We are the first to use CMA as an early fusion mechanism in HSD for videos, whose output is subsequently concatenated with modality-specific representations in a late fusion step before final prediction.
- **Incorporating on-screen text as a standalone modality in CMA:** We are the first work to use on-screen text as a standalone channel in the context of CMA, showing that on-screen text attending to concatenated transcript, audio, and video frames features acts as a useful additional feature in the context of HSD.
- **CMA-based interactions between the modalities:** We perform an analysis of the interactions between different modalities in the CMA setup, by testing different key-query combinations. We show that on-screen text performs best when used as a query attending to other modalities.

We release MM-HSD as an open-source benchmark for video-based HSD to support and advance ongoing research in this area¹.

2 RELATED WORK

HSD. Early social media platforms primarily supported textual content, which led to the development of the first HSD methods in the textual domain (HSD-t). They relied on keyword-based approaches, using predefined dictionaries of offensive words [29, 61]. Although these methods could achieve high precision for explicit slurs, they suffered from low recall, since HS often depends on context [45]. Also, due to the volume of online content, manual moderation is unscalable. To address these limitations, machine learning methods such as Naïve Bayes classifiers [38] and Support Vector Machines (SVMs) [27] gained popularity [2]. Later, deep learning architectures like Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks further improved performance [21, 60]. Today, transformer-based models are the standard in HSD [8, 12, 28, 43, 48].

Multi-modal Integration Techniques in HSD. Multi-modal HSD integrates modalities such as images, audio, and metadata to improve detection performance [59]. Multiple modalities could be combined using different fusion strategies, such as, for example, concatenation [12, 40, 67], and CMA [40, 47, 71]. Although concatenation has an advantage in its implementation simplicity, it fails to capture inter-dependencies between the modalities [70]. CMA has gained relevance in HSD in videos (HSD-v) [40, 71], and related domains such as toxicity, sexism and condescending language detection [4, 47, 65], and in general video classification [11], since it enables one modality to selectively focus on the most informative features of another [55]. Its importance is highlighted in studies

Table 1: HSD in videos. KEY – T: text modality (speech transcript), A: audio modality (waveform), V: video modality (frames), O: on-screen text (text in the frame), M: video meta-data, n/a: not applicable, MoE: Mixture of Experts, MPL: Multimodal Projection Layer in Llama 3.2-vision, CMA: cross-modal attention, CONCAT: concatenation fusion.

Reference	Modality					Fusion
	T	A	V	O	M	
[1, 69]	•					n/a
[34]					•	n/a
[12, 40, 67]	•	•	•			CONCAT
[40, 47]	•	•	•			CMA
[41]	•	•	•		•	MoE
[66]	•		•			MPL
[71]	•	•	•	•		Bimodal CMA
MM-HSD (ours)	•	•	•	•		CMA and CONCAT

such as [46], where the phrase “they shot another hamster” is benign, but “they shot another monkey” can be hateful if “monkey” is interpreted as a racial slur. This distinction can only be made when attending to both textual and visual modalities.

Multiple modalities can be combined with late or early fusion. Early fusion includes a combination of modalities directly before being fed into the decision model [18, 63]. Late fusion involves individual modalities to undergo high-level feature extraction and are then merged before passing through a final classification layer [44, 64]. Early fusion focuses on fine-grained interactions between modalities, whereas late fusion focuses on coarse-grained interactions [54]. **HSD in Videos.** HS in videos can be propagated via different channels, including spoken language, audio signals, visual content, and on-screen text. HSD-v methods differ based on which modalities are used, and how these modalities are combined (see Table 1). The first studies on HSD-v did not fully leverage multi-modal approaches but analyzed video transcripts [1, 69] or metadata such as YouTube titles, descriptions, and comments [34], which makes them essentially text-based approaches.

Recent methods combine transcript (T), audio (A), and video (V) using late fusion. Das et al. [12] showed that combining BERT, Vision Transformer (ViT), and MFCC improves Macro-F1 (M-F1) by 11.4% over their best unimodal setup. Wang et al. [67] found that late fusion also outperformed across English and Bilibili datasets. Maity et al. [47] introduced CMA for HSD-v in Hindi-English code-mixed videos, using LLaMA-3, Whisper, and VideoMAE. CMA was only applied in bimodal (transcript-audio or transcript-video) settings, omitting audio-video interactions. Video and audio modalities are encoded separately and independently attend to text tokens, generating “soft multi-modal tokens” that are passed back into the language model. This setup is mid- to late-fusion and remains strictly text-centric, without exploring alternative query modalities. Optical Character Recognition (OCR) of on-screen text (O) is omitted entirely.

Several models have been evaluated on the **HateMM** dataset [12], with their performance summarized in Table 2. HCC1 [40] sets the state-of-the-art M-F1 (0.848) and accuracy (0.854), combining HateXplain (T), CLIP (V), and CLAP (A) via late fusion. Authors [40] also explore a variant with CMA (MO-Hate) with an attention

¹<https://github.com/idiap/mmm-hsd>

Table 2: Performance comparison of models trained on HateMM [12]. Bold indicates the best performance and underline the second-best for each metric. We report our results over 5 runs as mean (std). Key — T: transcript, V: video, A: audio, O: on-screen text, M: macro average across both classes, H: hate class, F1: F1-score, ACC: unbiased accuracy, P: precision, R: recall.

Model	Architecture	T	V	A	O	ACC	M-F1	F1(H)	P(H)	R(H)	P(M)	R(M)
[12]	BERT, ViT, MFCC	•	•	•		.798	.790	.749	.742	.758	—	—
[40]	HXP, CLAP, CLIP	•	•	•		<u>.854</u>	<u>.848</u>	—	—	—	<u>.840</u>	<u>.800</u>
[66]	LLaMA-3.2-11B	•	•			.820	.820	.800	.800	.790	—	—
[71]	BERT, ViT, wav2vec + OCR + CMA	•	•	•	•	.849	.840	.876	.857	.896	—	—
MM-HSD (ours)	Detoxify, ViT, wav2vec, OCR + CMA	•	•	•	•	.878 (.009)	.874 (.009)	<u>.853 (.009)</u>	<u>.849 (.017)</u>	<u>.857 (.000)</u>	.874 (.010)	.875 (.008)

chain, where the transcript attends to audio (T→A), followed by text-audio attending to video ((T⊕A)→V). This chain underperforms in comparison to a simple concatenation, leading the authors to discard CMA. Moreover, it excludes on-screen text and only models bimodal interactions. The authors also report that Whisper significantly improves performance – a component we adopt as well. TCE-DBF [71] achieves the highest Micro-F1 (hate) score of 0.876, using CMA to combine text (transcript and on-screen), audio, and video – making it, alongside our work, the only model to incorporate on-screen text. The transcript and on-screen text are jointly processed using BERT without explicit modality signals, which may blur their semantic roles. In addition to this, modality interactions are limited: each modality is independently encoded (no early fusion), and text acts as the sole query in two distinct late-fusion cross-attention blocks (T→A and T→V). Also, alternative query-key combinations are not explored.

Vid+RM-FT [66] fine-tunes LLaMA-3.2-11B and LLaVA-Next-Video-7B on HateMM and a re-annotated version of the Hateful Memes dataset [37]. However, it excludes on-screen text and relies on a general-purpose backbone not optimized for HSD-t, increasing the risk of overfitting due to the dataset’s limited size. Lang et al. [41] introduce a Mixture-of-Experts (MoE) model, where modality-specific experts are combined with weights predicted by a router conditioned on all modalities. However, the fusion mechanism may be too simplistic to fully capture cross-modal interactions. Additionally, as in earlier work, BERT is fine-tuned and OCR is omitted – both limitations that may impact performance.

Prior CMA-based approaches use a single configuration without extensive comparison between possible setups. We are the first to evaluate CMA for HSD with early vs late fusion, query-key combinations, and modality integration/exclusion, which provides guidance for CMA design. Previous work [40, 47, 71] all rely on late fusion attention, and use text exclusively as query, omitting exploration of alternative modalities’ interactions. Only TCE-DBF [71] includes OCR by merging its output with the transcript. Koushik, Kanojia, and Treharne [40] attempted sequential CMA, but achieved poor performance, which made them discard CMA.

Finally, HS in videos datasets [12, 67] include frame-level annotations, but all their associated models [40, 66, 71] are video-level. A (partial) exception is [41] that segments videos for relevance scoring, but outputs a video-level prediction. Related unimodal work includes audio-level localisation via TTS [36] and word-level localisation [3]. No multi-modal HS models explore true frame-level localisation [41, 47].

3 DATASET

Contrasting with the abundance of textual HS datasets [8, 14, 15, 48], video-based datasets remain underdeveloped [42, 67]. Datasets that claim to address video-based hate-speech detection are often unimodal, containing only text comments or video transcriptions [1, 16, 26]. Many of these datasets are not publicly available [58, 69]², or merely provide links to video platforms [67]. Other datasets focus on related yet distinct concepts such as cyberbullying [19], sexism [4], or condescending language [65], and are therefore not wholly applicable to HS detection. Additionally, the applicability of some datasets is affected by the language in which HS is expressed: some datasets are in Hinglish (a mixture of English and Hindi) [47], or in Bengali [32], and therefore might propagate HS differently to datasets in English.

In this study, we use the HateMM dataset, a publicly available dataset comprising 1083 labeled videos from the BitChute platform [12], totalling 43 hours. This dataset contains 431 videos (39.8%) labeled as hate and 652 (60.2%) as non-hate. Video lengths vary significantly, ranging from a few seconds to over an hour, with an average duration of 2.40 minutes–2.56 minutes for hate videos and 2.28 minutes for non-hate videos. The dataset stands out for its diversity in how HS is represented, including videos where HS is conveyed through spoken words, displayed in text on video frames, or implied through actions depicted in the footage. Such a range is essential to underscore the potential of multi-modal models. In Figure 1, we present different examples of hate and non-hate videos from the dataset. Example (a) displays a Ku Klux Klan (Christian extremist, white supremacist, far-right hate group) gathering with a voiceover stating "We must segregate". While the image alone might have documentary purposes, the combination of visual and audio elements is unambiguously hateful. In contrast, example (b) features on-screen text explicitly calling to antisemitism, constituting HS through the text in the visual modality. This demonstrates the importance of including OCR of video frames into the model, enabling the detection of HS in on-screen text that might be missed by image-only analysis. On the other hand, examples (c) and (d) show images that could be interpreted as HS if there was an accompanying hateful text or voiceover. However, (c) is a clip from a mixed martial arts fight and (d) features news reports; both serve a purely informative purpose.

Additionally, the dataset comprises a variety of video types, which enhances its applicability and generalizability to real-world scenarios. Some videos consist of static images overlaid with text,

²We did not receive a reply from the authors to get access to the datasets.



Figure 1: Sample frames extracted from videos classified as **HATE ((a), (b)) and from videos classified as **NO HATE** ((c), (d)). Frames adapted from HateMM dataset [13], CC BY 4.0.**

accompanied by voiceovers or background music, while others include dynamic footage of individuals in public or private spaces.

4 METHODOLOGY

This section presents our approach for developing the MM-HSD multi-modal model. We begin by discussing the CMA mechanism [11] and how we propose to integrate CMA as an extra modality. Assuming the availability of pre-extracted embeddings for each modality (video, audio, transcript, and on-screen text), we then describe how CMA is integrated into the overall multi-modal architecture. This is followed by a comprehensive explanation of the embedding extraction process. Since the availability of video datasets for HSD is currently very limited, it is inefficient to fine-tune unimodal feature extractors (such as BERT) based on these small multi-modal datasets, as done in previous works [65, 67]. For this reason, we use already fine-tuned HS recognition models — such as Detoxify [28] — as text feature extractors, since they have been trained on much larger datasets for HSD-t. A detailed diagram of the complete pipeline is presented in Figure 2.

4.1 Models

CMA. CMA integrates information from different modalities by allowing the model to focus on specific aspects of one modality that are most relevant to another [11]. We explore the contribution of CMA by primarily evaluating it as an additional processing module output of which is fused with the modality-specific model outputs (see Figure 2, II) – this constitutes our MM-HSD model. We additionally evaluate its performance in two more distinct roles, which are illustrated in Figure 2 with color-coded paths: I) a late fusion layer that integrates the outputs of modality-specific models for video, audio, transcript and on-screen text, which we call CMA-LF, and III) a standalone feature extractor that directly processes

the raw embeddings from all modalities without using separate models for each modality, CMA-S. The goal is to further incorporate relationships between different modalities [20].

CMA consists of an attention mechanism with Query (Q), Key (K), and Value (V). It is expressed as

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

where d_k is the dimensionality of the key vectors. In this work, Q represents the modalities that extract relevant information from other modalities. K are the modalities to which the query attends to identify the most relevant information. V holds the actual data, from which the most relevant parts are extracted based on the query-key interaction. The choice of Q, K, and V values is studied, and thus different combinations are tuned to find the optimal combination. The K and V modalities are identical, while the Q is assigned to the remaining modality that is not used for K and V. For instance, in a transcript-video-audio setup, some of the scenarios in which the model is tested include:

- $K=[\text{audio}, \text{transcript}], Q=[\text{video}], V=[\text{audio}, \text{transcript}]$
- $K=[\text{video}, \text{transcript}], Q=[\text{audio}], V=[\text{video}, \text{transcript}]$

...

Note that in the case of multiple modalities serving as the key, they are concatenated along the sequence dimension. This means that the embeddings from different modalities are stacked sequentially, allowing the attention mechanism to process them as a unified extended sequence rather than treating each modality separately [40].

Model Ensembling with CMA as an extra modality (MM-HSD). The pipeline for MM-HSD is described in Figure 2, II. CMA is applied directly to the raw modality embeddings, and its output is concatenated with the outputs of the individual modality encoders before a final classification head. We use concatenated features of T, V, and A modalities as keys, keeping O as a query, a combination that was experimentally found to yield the best performance (see Appendix A, Table 7). An analysis of the selection process is presented in Section 5.5. We report the performance of MM-HSD in Table 2.

CMA as a standalone feature extractor (CMA-S). We apply CMA directly to the raw modality embeddings. We evaluate the predictive power of CMA in the early fusion settings (see Figure 2, III). The output of CMA is then passed through a feedforward layer for final classification, with no additional unimodal transformations. We use concatenated features of T, V, and A modalities as keys, and O as a query.

CMA as a late fusion strategy (CMA-LF). CMA is used as a late fusion strategy, on the outputs of the individual modality encoders before the final classification head (see Figure 2, I).

4.2 Embeddings extraction and encoding

We now describe the pre-processing and feature extraction steps for each individual modality, followed by the subsequent encoding stage.

Video. Video files are sampled at one frame-per-second, with a maximum of 100 frames extracted per video, following [12, 40]. Shorter videos are padded with blank frames to reach 100 frames.

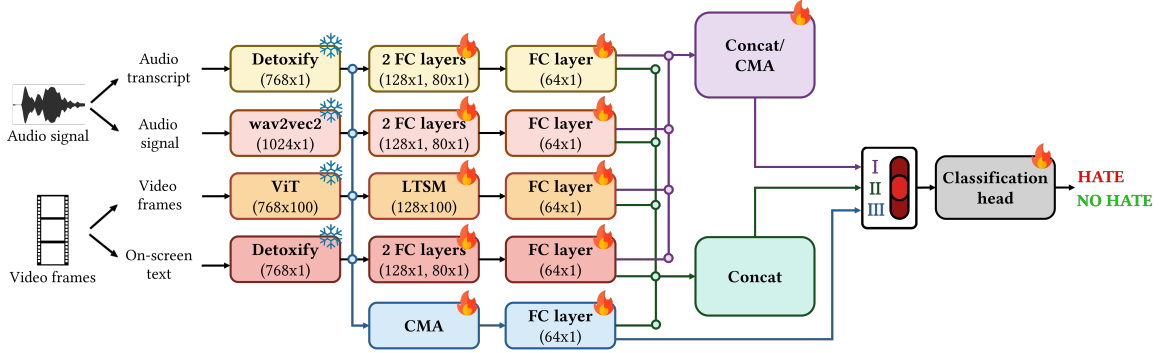


Figure 2: Our tetra-modal HSD architecture. The four modalities are audio transcripts, audio signal, video frames, and on-screen text. The embeddings for each modality are obtained by pre-trained models and then used in three different experimental setups, controlled by a switch module with three configurations (I, II, III). **I) Late Fusion Experiments :** Each modality is first independently encoded (e.g. by LSTM or Fully Connected (FC) layers), and their outputs are fused using either concatenation or CMA. **II) Late Fusion with CMA as Additional Modality :** CMA is applied to raw modality embeddings to generate an additional feature, which is then combined (via concatenation) with the encoded features. **III) Early Fusion with CMA as Unique Feature Extractor :** Raw modality embeddings are fused directly using CMA, without any modality-specific processing. The final classification head outputs a binary label: **NO HATE** or **HATE**.

The embeddings are generated using a pretrained ViT [17]. We apply the ViT to each frame independently and concatenate their embeddings. ViT has been chosen for visual feature extraction due to its ability to model long-range dependencies and thus capture global context through self-attention mechanisms – unlike convolutional networks, which rely on local receptive fields. In addition, we aim to separate the concerns of image context contribution and on-screen text cues. Newer models like CLIP [56] have shown to have higher sensitivity towards on-screen text than to the image components, such as texture [51], while ViT should be able to focus on the context within the image rather than on the on-screen text. **Audio and Transcriptions.** From each video, we extract its audio, which we then transcribe. This way we separate the content of the human speech, which is contained in the transcript and could be processed by a pre-trained hate-speech encoder, from the emotional cues contained in the speech, such as aggressiveness, which are contained in audio [12]. To obtain acoustic features, we use a pre-trained wav2vec2-large-xlsr-53 model [5, 25], a self-supervised speech recognition model composed of a convolutional feature encoder and a transformer network fine-tuned for the English language. Finally, we transcribe the audio using OpenAI’s Whisper model [57]. This text is then encoded into low-dimensional vectors using Detoxify [28], a model based on RoBERTa [43] and specifically trained to detect various forms of toxic and hateful language. A complementary ablation on removing stopwords on audio transcripts is included in Appendix B. Stopwords removal led to a performance drop and was not used in the main experiments.

On-screen text. To extract the printed text, captions, or other textual elements within the video frames, we sample one frame per second. The text from each frame is extracted using PaddleOCR [52], since it can handle different text orientations and noisy frames with minimal postprocessing. Once extracted, text coherence is enhanced

by applying the following postprocessing steps. First, the text is cleaned by removing any unwanted characters, retaining only alphanumeric symbols, common punctuation, and apostrophes for contractions. Next, de-duplication is performed to reduce redundancy introduced by frames extracted at short intervals. Duplicate or highly similar text segments are filtered out by comparing their similarity scores and discarding those over 90% resemblance to an already retained entry. Finally, overlapping text fragments are merged by identifying the longest matching sequences and combining them into a cohesive output, ensuring a more accurate and readable final text. Similarly to the audio transcriptions, we embed the OCR output with a Detoxify model.

We tested the effect of replacing MM-HSD’s original feature extractors one modality at a time – using BERT instead of Detoxify for transcript and on-screen text, MFCC instead of wav2vec for audio, and InceptionV3 instead of ViT for video. These alternative extractors underperformed compared to the chosen ones. For instance, MM-HSD outperforms all alternatives by 2.4–4.7% in M-F1. This supports the chosen extractors for each modality.

Encoding of individual modalities. This component is used in setups I and II. The video modality is represented by 768-dimensional features extracted from each frame, which are fed into an LSTM network, followed by an FC layer for hate video classification. LSTM networks excel at detecting patterns and dependencies over time, such as actions within the video stream, making it a suitable choice for understanding the dynamics between frames. For text, we input the 768-dimensional embeddings from both the audio and OCR transcripts into separate neural networks, each consisting of three FC layers. For the audio signal, the 1024-dimensional wav2vec2 features serve as input for training a three-FC layer neural network.

Table 3: Results for unimodal experiments models, CMA as a standalone feature extractor (CMA-S), CMA as an extra modality (MM-HSD), CMA as late fusion (CMA-LF) and modality-specific models fused with concatenation without CMA (w/o CMA). We report our results over 5 runs as mean (std). Key – T: transcript, V: video, A: audio, O: on-screen text, M: macro average across both classes, H: hate class, F1: F1-score, ACC: unbiased accuracy, P: precision, R: recall.

Model	ACC	M-F1	F1(H)	P(H)	R(H)
T	.820 (.012)	.816 (.012)	.790 (.012)	.765 (.019)	.816 (.009)
O	.636 (.014)	.594 (.011)	.464 (.012)	.596 (.032)	.381 (.016)
A	.784 (.019)	.778 (.018)	.742 (.018)	.739 (.039)	.746 (.030)
V	.761 (.027)	.751 (.024)	.702 (.020)	.730 (.055)	.679 (.017)
CMA-S [†]	.850 (.006)	.846 (.006)	.820 (.006)	.818 (.016)	.821 (.008)
MM-HSD [†]	.878 (.009)	.874 (.009)	.853 (.009)	.849 (.017)	.857 (.000)
w/o CMA	.846 (.013)	.842 (.014)	.817 (.019)	.805 (.028)	.832 (.052)
CMA-LF [†]	.842 (.024)	.837 (.024)	.810 (.028)	.812 (.057)	.813 (.057)

[†] Query: On-screen text. Key: concatenation of transcripts, video, and audio.

5 VALIDATION

5.1 Experimental setup

Given a limited amount of data, we adopt a 5-fold cross-validation strategy on 85% of the data, reserving the remaining 15% for testing. Each fold includes 698 training and 175 validation samples, with a batch size of 8. All runs are executed on an NVIDIA RTX 3090. We employ elastic net regularization, combining L1 and L2 penalties to enhance sparsity and stability. A ReduceLROnPlateau scheduler reduces the learning rate by a factor of 0.1 when validation loss does not decrease for 6 consecutive epochs. Early stopping is triggered if no improvement is observed over a patience period. Dropout is applied after each FC layer to further regularize the model. To account for class imbalance, we use a weighted cross-entropy loss. Hyperparameter tuning is performed over a fixed grid, with the learning rate $\eta \in \{10^{-3}, 10^{-4}, 10^{-5}\}$, L1 penalty $\lambda_1 \in \{10^{-3}, 10^{-4}, 10^{-5}\}$, L2 penalty $\lambda_2 \in \{10^{-4}, 10^{-5}, 10^{-6}\}$, dropout rate $d \in \{0.3, 0.4, 0.5\}$, and early stopping patience $p \in \{5, 10\}$.

Metrics. Model performance is assessed using both overall and class-specific metrics. Unbiased Accuracy (ACC) – the mean of class-wise recall in binary classification – ensures performance is not skewed by class imbalance. Hate F1 (F1(H)) captures performance specifically on the Hate class, while M-F1 averages F1 scores across both Hate and No Hate classes. We additionally report Hate Precision (P(H)) and Hate Recall (R(H)). The best-performing models are chosen according to the training and validation loss trends and the corresponding validation F1, recall, and precision macro-averaged and hate-specific scores. In Table 2, we summarize the comparative performance of state-of-the-art models on HateMM [12] dataset. Note that while TCE-DBF [71] reports higher hate metrics, our model outperforms in macro metrics. Additionally, results of [71] are from a single run with unknown variability.

5.2 Baseline models comparison

We evaluate the predictive strength of the individual modalities of our model. In the first rows of Table 3 we present the unimodal baselines (T, V, A, O), and the CMA as a standalone feature extractor

Table 4: Results for CMA as an extra modality (MM-HSD). We report our results over 5 runs as mean (std). Key – Mod: modalities, K: key, Q: query, T: transcript, O: on-screen text, A: audio, V: video, M: macro average across both classes, H: hate class, F1: F1-score, ACC: unbiased accuracy, P: precision, R: recall.

Mod.	K	Q	ACC	M-F1	F1(H)	P(H)	R(H)
TO	T	O	.830 (.006)	.825 (.005)	.796 (.006)	.793 (.014)	.800 (.014)
TA	T	A	.828 (.025)	.823 (.024)	.796 (.023)	.786 (.046)	.806 (.007)
TV	T	V	.841 (.006)	.837 (.006)	.811 (.006)	.799 (.009)	.822 (.007)
OA	A	O	.805 (.028)	.801 (.027)	.774 (.024)	.749 (.048)	.803 (.014)
OV	V	O	.775 (.009)	.768 (.010)	.730 (.014)	.726 (.005)	.733 (.026)
AV	A	V	.808 (.026)	.799 (.030)	.759 (.041)	.788 (.023)	.733 (.065)
TOA	TO	A	.834 (.011)	.830 (.011)	.805 (.014)	.787 (.023)	.825 (.037)
TOV	TV	O	.838 (.019)	.834 (.019)	.807 (.022)	.800 (.035)	.816 (.040)
TVA	TV	A	.849 (.007)	.845 (.007)	.819 (.009)	.811 (.014)	.829 (.021)
OAV	OA	V	.821 (.023)	.815 (.026)	.781 (.035)	.789 (.013)	.775 (.059)
TOAV	TAV	O	.878 (.009)	.874 (.009)	.853 (.009)	.849 (.017)	.857 (.000)

baseline (CMA-S), which processes the four modalities in the early fusion setting, as described in Section 4.1. Consistent with the literature [12, 47], we observe that the T modality is the most effective unimodal model, achieving an M-F1 score of 0.816. In contrast, the O modality yields the lowest performance, with an M-F1 score of 0.594. This disparity is expected, as not all videos in the dataset contain on-screen text, making the O modality ineffective in some instances. CMA-S uses the O modality as the query and the T, A, and V modalities as keys. The performance of CMA-S being higher (M-F1 of 0.846) than unimodal baselines is to be expected, as it already constitutes a fully multi-modal model that combines information from all modalities. As such, it already reflects the benefits of multi-modal HSD and the advantages of using CMA.

We also compare the CMA mechanism used independently as a feature extractor, CMA-S, with CMA incorporated as an additional modality within the full multi-modal framework, MM-HSD. As shown in Table 3, combining CMA and the individual models results in a boost in performance, increasing M-F1 score from 0.846 to 0.874. The reason behind this is that leveraging modality-specific encoders allows the model to use both cross-modal interaction and the specialized representations captured by each unimodal encoder. Note that we maintain the query-key pair aforementioned, with O as query and TVA as key for both CMA-S and MM-HSD. For completeness, we evaluate CMA as a late fusion strategy for the modality-specific models, CMA-LF, which underperforms (M-F1 score of 0.837) compared to MM-HSD. We use MM-HSD for the rest of the experiments.

5.3 Decreasing the number of modalities

In Table 4, we show how model performance changes when individual modalities (T, O, A, V) are removed from the full tetra-modal setup of MM-HSD. The key-query pairs used are chosen according to the training and validation metrics obtained when using CMA-S. When dropping any single modality from the tetra-modal configuration (M-F1 = 0.874), performance declines to between 0.815 and 0.845, already showing that all modalities contribute unique information for HSD. The largest drop occurs when the T modality is

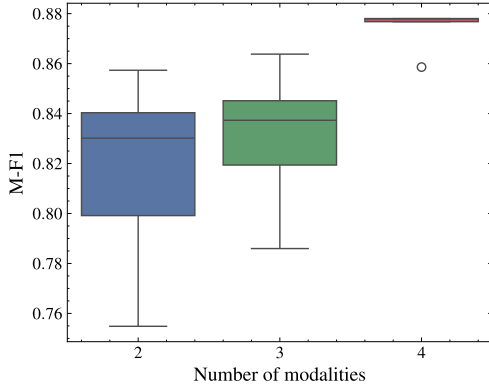


Figure 3: Impact of modality combinations on the M-F1 score. Each box represents the distribution of mean M-F1 scores (mean computed over five runs with different random seeds) for each specific combination of modalities, grouped by the number of modalities used — from bimodal to tetra-modal setups. The combinations include all possible subsets of the four modalities (T, A, V, and O) with key-query pairs selected on training and validation metrics on CMA-S. The trend shows that M-F1 generally increases as more modalities are integrated.

removed, and the lowest when the O modality is removed. Likewise, in the transition from a trimodal system (best M-F1 = 0.845) down to any bimodal setup, scores drop further, between 0.768 and 0.837. The same modalities that drive performance in the tetra-modal setup also contribute the most when fewer modalities are used. Reducing the number of input modalities increases the range of the mean M-F1 scores, which is illustrated in Figure 3. We further observe that the specific combination of modalities used as keys becomes more influential — leading to substantial variation within 2-, 3-, and 4- modality settings.

To further isolate the contribution of each modality within the attention mechanism, we perform an additional experiment on MM-HSD in which modalities are selectively excluded from the CMA block only, with the O modality fixed as the query. The resulting attention output is then concatenated with the individual modality encoder outputs before classification, following the original architecture. We observe consistent drops in M-F1 and F1(H) when any modality is excluded (see Appendix C), further confirming the need to include all modalities.

5.4 Removing CMA features as a modality

To assess the contribution of the CMA mechanism to overall model performance, we evaluate the configuration with only the modality-specific encoders, and excluding CMA as a feature extractor. This enables a clearer understanding of the added value provided by CMA beyond the capabilities of the individual modalities in isolation. For this comparison, we leverage concatenation as the fusion technique. We observe that incorporating CMA as an extra modality results in a performance boost: M-F1 score of 0.878 (MM-HSD) against 0.846 (w/o CMA), as shown in Table 3. Moreover, the incorporation of CMA leads to a significant reduction in standard

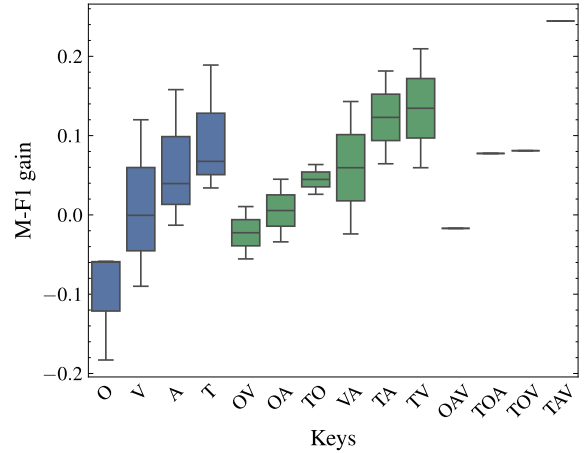


Figure 4: The dependence between keys used in CMA fusion and the average M-F1 performance gain in comparison to a corresponding unimodal baseline (model trained on the query modality). For each key, the distribution includes all possible query options that do not share any modalities with that key. The distribution is based on the average M-F1 gain for CMA-S and CMA-LF. Key – T: transcript, V: video, A: audio, O: on-screen text, M-F1: Macro-F1 score.

deviation, indicating improved consistency and robustness across evaluation runs.

5.5 Performance trends by CMA configuration

We evaluate the relationship between different modalities using CMA as a fusion strategy in two experimental setups: CMA-S and CMA-LF. In these configurations, the difference lies in the input to CMA-S being raw embeddings, and the input to CMA-LF being encoded embeddings. In both configurations, we explore all possible query-key combinations where one modality (the query) attends to a set of other modalities (the keys), which are concatenated together. Conceptually, when a modality A attends to $\text{concat}(B_1, \dots, B_n)$, this reflects the amount of useful information modality A can extract from the modalities it attends to. The full list of key-query pairs used in these experiments is provided in Appendix A, Tables 6 and 7. Note that these experiments were conducted using a single-seed setup.

To evaluate the impact of these interactions, we compare the M-F1 scores of CMA-S and CMA-LF against those of the unimodal baselines, where the modality used in the unimodal model matches the query used in CMA. The gain in performance is then calculated as the difference in M-F1 between the CMA-based model and its unimodal counterpart. First, we validate the stability of the results by examining the correlation between performance gains for both CMA-S and CMA-LF. We observe a very strong positive correlation of 91%. This suggests that the patterns observed hold for both late and early fusion, and therefore can be generalized. To increase robustness, we take the average between the gain in performance for late and early setups, and conduct our analysis based on these averaged values (see Figure 4). We note that in general, whenever O is used as a part of a key, the performance drops, meaning that

Table 5: Efficiency metrics for unimodal and multimodal models. Key – CMA used as standalone (CMA-S), additional modality (MM-HSD), late fusion (CMA-LF), and removed (w/o CMA), TTE: Train Time per Epoch, TTT: Total Train Time, TT: Test Time, Par: Parameters.

Model	TTE (s)	TTT (s)	TT (s)	# Par (M)	Size (MB)
A	0.540	73.162	0.046	0.147	0.562
T	0.441	65.818	0.058	0.123	0.470
O	0.426	18.075	0.050	0.123	0.470
V	0.462	31.427	0.041	1.279	4.880
CMA-S	1.124	155.917	0.068	2.953	11.266
MM-HSD	1.465	293.022	0.060	4.626	17.648
w/o CMA	0.975	70.013	0.065	1.673	6.381
CMA-LF	1.271	81.223	0.089	1.722	6.570

on-screen text does not allow enrichment of other modalities when used as a key. Using a transcript as a part of the key, allows the queries to gain a boost in performance, meaning that stronger modalities can often enrich the weaker modalities when used as a key. We note that the greatest gain happens when O is attending to T, A, and V.

5.6 Computational cost analysis

We summarize the computational cost of our models in Table 5. Training time per epoch is measured on 698 training and 175 validation samples, and the results are averaged over 5 folds. There is variation in training time across models introduced by early stopping. We observe that models with limited training capacity (e.g. unimodal O baseline) complete training sooner. However, CMA runtime scales quadratically with input size, resulting in much longer training times for models such as CMA-S and MM-HSD. Inference time is calculated over 155 test samples, and remains consistently low across all models, even though using CMA greatly increases parameter count in models such as CMA-S and MM-HSD.

Note that feature extraction using Detoxify, ViT, and wav2vec2 is performed offline, and not counted towards the total runtime of the different models. These feature extractors are large (Detoxify contains 109M parameters, ViT – 86M, and wav2vec2 – 315M) and not well-suited for on-device deployment. However, the final classifiers are relatively lightweight, with a maximum size of 4.6M parameters for MM-HSD, making on-device use possible if the extracted features are available.

6 CONCLUSION

We proposed MM-HSD, a model for multi-modal HSD in videos, and analyzed the contribution of audio, video, speech transcripts, and on-screen text modalities. We showed that incorporating CMA as an additional modality alongside modality-specific encoders leads to improved performance compared to both unimodal models and late-fusion strategies. Our results suggest that decoupling visual feature extraction from on-screen text recognition – by using a general-purpose image encoder (e.g. ViT) in combination with a separate OCR module – allows for a more nuanced HSD. By analyzing query-key pairs in the context of CMA, we showed that on-screen text attending to transcripts, audio, and video yields the best performance.

As future work, one could evaluate whether OCR-to-speech conversion improves HS classification within the CMA framework [6, 7]. Furthermore, it would be interesting to explore frame-level localization using temporal CMA [49], which may improve model explainability by showing which frames contributed to the final classification. Lastly, MM-HSD has been trained to optimize hate classification on HateMM solely. Future research should focus on generalization and validation across additional datasets.

ACKNOWLEDGMENTS

P.R.-R. received the support of a fellowship from "la Caixa" Foundation (ID 100010434), with fellowship code LCF/BQ/EU23/12010085. C.C.-C. received support of a fellowship from "Mutua Madrileña Foundation".

REFERENCES

- [1] Cleber Alcântara, Viviane Moreira, and Diego Feijo. 2020. Offensive Video Detection: Dataset and Baseline Results. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*. 4309–4319.
- [2] Ahlam Alrehili. 2019. Automatic Hate Speech Detection on Social Media: A Brief Survey. In *Proceedings of the 2019 IEEE/ACS 16th International Conference on Computer Systems and Applications (AICCSA)*. IEEE, 1–6.
- [3] Jimmyeong An, Wonjun Lee, Yejin Jeon, Jungseul Ok, Yunsu Kim, and Gary G. Lee. 2024. An Investigation into Explainable Audio Hate Speech Detection. In *Proceedings of the 25th Annual Meeting of the Special Interest Group on Discourse and Dialogue*. Association for Computational Linguistics, Kyoto, Japan, 533–543. doi:10.18653/v1/2024.sigdial-1.45
- [4] Iván Arcos and Paolo Rosso. 2024. Sexism Identification on TikTok: A Multimodal AI Approach with Text, Audio, and Video. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction: 15th International Conference of the CLEF Association, CLEF 2024, Grenoble, France, September 9–12, 2024, Proceedings, Part I* (Grenoble, France). Springer-Verlag, Berlin, Heidelberg, 61–73. doi:10.1007/978-3-031-71736-9_2
- [5] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 12449–12460.
- [6] Kirtilekha Bhesra and Akshay Agarwal. 2025. A Multi-modal Framework to Counter Hate Speeches. In *Pattern Recognition*, Apostolos Antonacopoulos, Subhasis Chaudhuri, Rama Chellappa, Cheng-Lin Liu, Saumik Bhattacharya, and Umapada Pal (Eds.). Springer Nature Switzerland, Cham, 197–207.
- [7] Kirtilekha Bhesra, Shivam A. Shukla, and Akshay Agarwal. 2024. Audio vs. Text: Identify a Powerful Modality for Effective Hate Speech Detection. In *The Second Tiny Papers Track at ICLR 2024*.
- [8] Tommaso Caselli, Valerio Basile, Jelena Mitrović, and Michael Granitzer. 2021. HateBERT: Retraining BERT for Abusive Language Detection in English. In *Proceedings of the 5th Workshop on Online Abuse and Harms (WOAH 2021)*, Aida Mostafazadeh Davani, Douwe Kiela, Mathias Lambert, Bertie Vidgen, Vinodkumar Prabhakaran, and Zeerak Waseem (Eds.). Association for Computational Linguistics, Online, 17–25. doi:10.18653/v1/2021.woah-1.3
- [9] Lu Cheng, Kai Shu, Siqi Wu, Yasin N. Silva, Deborah L. Hall, and Huan Liu. 2020. Unsupervised Cyberbullying Detection via Time-Informed Gaussian Mixture Model. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (Virtual Event, Ireland) (Cikm '20)*. Association for Computing Machinery, New York, NY, USA, 185–194. doi:10.1145/3340531.3411934
- [10] Anusha Chhabra and Dinesh K. Vishwakarma. 2023. A Literature Survey on Multimodal and Multilingual Automatic Hate Speech Identification. *Multimedia Systems* 29, 3 (2023), 1203–1230.
- [11] Lu Chi, Guiyu Tian, Yadong Mu, and Qi Tian. 2019. Two-Stream Video Classification with Cross-Modality Attention. In *Proceedings of the IEEE/CVF international conference on computer vision workshops*.
- [12] Mithun Das, Rohit Raj, Punyajoy Saha, Binny Mathew, Manish Gupta, and Animesh Mukherjee. 2023. HateMM: A Multi-Modal Dataset for Hate Video Classification. *Proceedings of the International AAAI Conference on Web and Social Media* 17, 1 (Jun. 2023), 1014–1023. doi:10.1609/icwsm.v17i1.22209
- [13] Mithun Das, Rohit Raj, Punyajoy Saha, Binny Mathew, Manish Gupta, and Animesh Mukherjee. 2023. HateMM: A Multi-modal Dataset for Hate Video Classification. doi:10.5281/zenodo.7799469

- [14] Thomas Davidson, Debasmita Bhattacharya, and Ingmar Weber. 2019. Racial Bias in Hate Speech and Abusive Language Detection Datasets. In *Proceedings of the Third Workshop on Abusive Language Online*, Sarah T. Roberts, Joel Tetreault, Vinodkumar Prabhakaran, and Zeerak Waseem (Eds.). Association for Computational Linguistics, Florence, Italy, 25–35. doi:10.18653/v1/W19-3504
- [15] Ona de Gibert, Naiara Perez, Aitor García-Pablos, and Montse Cuadros. 2018. Hate Speech Dataset from a White Supremacy Forum. In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*, Darja Fišer, Ruihong Huang, Vinodkumar Prabhakaran, Rob Voigt, Zeerak Waseem, and Jacqueline Wernimont (Eds.). Association for Computational Linguistics, Brussels, Belgium, 11–20. doi:10.18653/v1/W18-5102
- [16] Abreham G. Debele and Michael M. Woldeyohannis. 2022. Multimodal Amharic Hate Speech Detection Using Deep Learning. In *2022 International Conference on Information and Communication Technology for Development for Africa (ICT4DA)*. IEEE, 102–107.
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xi-aohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=YicbFdNTTy>
- [18] Chi T. Duong, Remi Lebre, and Karl Aberer. 2017. Multimodal Classification for Analysing Social Media. arXiv:1708.02099 [cs.CL]
- [19] Eniafe Festus Ayetiran and Özlem Özgöbek. 2024. A Review of Deep Learning Techniques for Multimodal Fake News and Harmful Languages Detection. *IEEE Access* 12 (2024), 76133–76153. doi:10.1109/access.2024.3406258
- [20] Jinmiao Fu, Shaoyuan Xu, Huidong Liu, Yang Liu, Ning Xie, Chien-Chih Wang, Jia Liu, Yi Sun, and Bryan Wang. 2022. CMA-CLIP: Cross-Modality Attention Clip for Text-Image Classification. In *2022 IEEE International Conference on Image Processing (ICIP)*. 2846–2850. doi:10.1109/icip46576.2022.9897323
- [21] Ankita Gandhi, Param Ahir, Kinjal Adhvaryu, Pooja Shah, Ritika Lohiya, Erik Cambria, Soujanya Poria, and Amir Hussain. 2024. Hate Speech Detection: A Comprehensive Review of Recent Works. *Expert Systems* (2024).
- [22] Raul Gomez, Jaume Gibert, Lluís Gomez, and Dimosthenis Karatzas. 2020. Exploring Hate Speech Detection in Multimodal Publications. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. 1470–1478.
- [23] Vaishali U. Gongane, Mousami V. Munot, and Alwin D. Anuse. 2022. Detection and Moderation of Detrimental Content on Social Media Platforms: Current Status and Future Directions. *Social Network Analysis and Mining* 12, 1 (2022), 129.
- [24] Satya K. Gorti, Noël Vouitsis, Junwei Ma, Keyvan Golestani, Maksims Volkovs, Animesh Garg, and Guangwei Yu. 2022. X-Pool: Cross-Modal Language-Video Attention for Text-Video Retrieval. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5006–5015.
- [25] Jonas Grosman. 2021. Fine-tuned XLSR-53 Large Model for Speech Recognition in English. <https://huggingface.co/jonatasgrosman/wav2vec2-large-xlsr-53-english>.
- [26] Shrey Gupta, Pratyush Priyadarshi, and Manish Gupta. 2023. Hateful Comment Detection and Hate Target Type Prediction for Video Comments. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. Acm, Birmingham United Kingdom, 3923–3927. doi:10.1145/3583780.3615260
- [27] Karimah M. Hana, Said Al Faraby, and Arif Bramantoro. 2020. Multi-Label Classification of Indonesian Hate Speech on Twitter Using Support Vector Machines. In *2020 International Conference on Data Science and Its Applications (ICoDSA)*. IEEE, 1–7.
- [28] Laura Hanu and Unitary team. 2020. Detoxify. Github. <https://github.com/unitaryai/detoxify>.
- [29] Hatebase. 2025. Hatebase is a collaborative, regionalized repository of multilingual hate speech. <https://hatebase.org> Accessed: 2025-02-27.
- [30] Ming S. Hee, Shivam Sharma, Rui Cao, Palash Nandi, Preslav Nakov, Tanmoy Chakraborty, and Roy Ka-Wei Lee. 2024. Recent Advances in Online Hate Speech Moderation: Multimodality and the Role of Large Models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 4407–4419. doi:10.18653/v1/2024.findings-emnlp.254
- [31] Eftekar Hossain, Omar Sharif, Mohammed M. Hoque, M. A. Akber Dewan, Nazmul Siddique, and Md. A. Hossain. 2022. Identification of Multilingual Offense and Troll from Social Media Memes Using Weighted Ensemble of Multimodal Features. *Journal of King Saud University - Computer and Information Sciences* 34, 9 (2022), 6605–6623. doi:10.1016/j.jksuci.2022.06.010
- [32] Mohd. I. Hossain, Faisal Hossain, and Rashedur M. Rahman. 2021. Bangla Hate Speech Detection in Videos Using Machine Learning. In *2021 IEEE 12th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON)*. IEEE, New York, NY, USA, 0347–0351. doi:10.1109/uemcon53757.2021.9666550
- [33] Olga Jubany and Malin Roiha. 2016. Backgrounds, Experiences and Responses to Online Hate Speech: A Comparative Cross-Country Analysis. *Online report*. Barcelona: University of Barcelona (2016).
- [34] Rajeshwari Kandakatla. 2016. *Identifying Offensive Videos on YouTube*. Master’s thesis. Wright State University. http://rave.ohiolink.edu/etdc/view?acc_num=wright1484751212961772 Available at OhioLINK Electronic Theses and Dissertations Center.
- [35] Simon Kemp. 2025. Digital 2025: The State of Social Media in 2025. (February 2025). <https://datareportal.com/reports/digital-2025-sub-section-state-of-social>
- [36] Hareem Kibriya, Ayesha Siddiq, Wazir Z. Khan, and Muhammad K. Khan. 2024. Towards Safer Online Communities: Deep Learning and Explainable AI for Hate Speech Detection and Classification. *Computers and Electrical Engineering* 116 (2024), 109153.
- [37] Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2020. The Hateful Memes Challenge: Detecting Hate Speech in Multimodal Memes. *Advances in neural information processing systems* 33 (2020), 2611–2624.
- [38] Kelvin K. Kiilu, George Okeyo, Richard Rimiru, and Kennedy Ogada. 2018. Using Naïve Bayes Algorithm in Detection of Hate Tweets. *International Journal of Scientific and Research Publications* 8, 3 (2018), 99–107.
- [39] Dirk Kindermann. 2023. Against ‘Hate Speech’. *Journal of Applied Philosophy* 40, 5 (2023), 813–835. doi:10.1111/japp.12648
- [40] Girish A. Koushik, Diptesh Kanojia, and Helen Treharne. 2025. Towards a Robust Framework for Multimodal Hate Detection: A Study on Video vs. Image-based Content. In *Companion Proceedings of the ACM on Web Conference 2025* (Sydney NSW, Australia) (WWW’25). Association for Computing Machinery, New York, NY, USA, 2014–2023. doi:10.1145/3701716.3718382
- [41] Jian Lang, Rongpei Hong, Jin Xu, Yili Li, Xovee Xu, and Fan Zhou. 2025. Biting off More Than You Can Detect: Retrieval-Augmented Multimodal Experts for Short Video Hate Detection. In *The Web Conference 2025*. <https://openreview.net/forum?id=GrJYmDzfW>
- [42] Phillip Lippe, Nithin Holla, Shantanu Chandra, Santhosh Rajamanickam, Georgios Antoniou, Ekaterina Shutova, and Helen Yannakoudakis. 2020. A Multimodal Framework for the Detection of Hateful Memes. arXiv:2012.12871 [cs.CL]
- [43] Yinhuan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. arXiv:1907.11692 [cs.CL]
- [44] Zhiyu Ma, Shaowen Yao, Liwen Wu, Song Gao, and Yunqi Zhang. 2022. Hateful Memes Detection Based on Multi-Task Learning. *Mathematics* 10, 23 (2022), 4525.
- [45] Sean MacAvaney, Hao-Ren Yao, Eugene Yang, Katina Russell, Nazli Goharian, and Ophir Frieder. 2019. Hate Speech Detection: Challenges and Solutions. *PLoS one* 14, 8 (2019).
- [46] Kosioschukwu J. Madukwe, Xiaoying Gao, and Bing Xue. 2022. Token Replacement-Based Data Augmentation Methods for Hate Speech Detection. *World Wide Web* 25, 3 (2022), 1129–1150.
- [47] Krishanu Maity, A. S. Poornash, Sriparna Saha, and Pushpak Bhattacharyya. 2024. ToxVidLM: A Multimodal Framework for Toxicity Detection in Code-Mixed Videos. In *Findings of the Association for Computational Linguistics: ACL 2024*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 11130–11142. doi:10.18653/v1/2024.findings-acl.663
- [48] Binny Mathew, Punyajoy Saha, Seid M. Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2021. HateXplain: A Benchmark Dataset for Explainable Hate Speech Detection. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 14867–14875.
- [49] Otniel-Bogdan Mercea, Thomas Hummel, A. Sophia Koepke, and Zeynep Akata. 2022. Temporal and Cross-modal Attention for Audio-Visual Zero-Shot Learning. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XX* (Tel Aviv, Israel). Springer-Verlag, Berlin, Heidelberg, 488–505. doi:10.1007/978-3-031-20044-1_28
- [50] Nanli S. Mullah and Wan M. N. W. Zainon. 2021. Advances in Machine Learning Algorithms for Hate Speech Detection in Social Media: A Review. *IEEE Access* 9 (2021), 88364–88376. doi:10.1109/access.2021.3089515
- [51] David A. Noever and Samantha E. M. Noever. 2021. Reading Isn’t Believing: Adversarial Attacks On Multi-Modal Neurons. arXiv:2103.10480 [cs.LG]
- [52] PaddlePaddle. 2025. PaddleOCR: Multi-Language OCR System. <https://github.com/PaddlePaddle/PaddleOCR>. Accessed: 2025-04-11.
- [53] Konstantinos Perifanos and Dionysis Goutsos. 2021. Multimodal Hate Speech Detection in Greek Social Media. *Multimodal Technologies and Interaction* 5, 7 (2021). doi:10.1039/mti5070034
- [54] Vittorio Pipoli, Federico Bolelli, Sara Sarto, Marcella Cornia, Lorenzo Baraldi, Costantino Grana, Rita Cucchiara, and Elisa Ficarra. 2025. Semantically Conditioned Prompts for Visual Recognition Under Missing Modality Scenarios. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. 4968–4977. doi:10.1109/WACV61041.2025.00486
- [55] R. G. Praveen and Jahangir Alam. 2024. Recursive Joint Cross-Modal Attention for Multimodal Fusion in Dimensional Emotion Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. 4803–4813.

- [56] Alec Radford, Jong W. Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models from Natural Language Supervision. In *International conference on machine learning*. PMLR, 8748–8763.
- [57] Alec Radford, Jong W. Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust Speech Recognition via Large-Scale Weak Supervision. In *International conference on machine learning*. PMLR, 28492–28518.
- [58] Aneri Rana and Sonali Jha. 2022. Emotion Based Hate Speech Detection using Multimodal Learning. arXiv:2202.06218 [cs.LG]
- [59] Anchal Rawat, Santosh Kumar, and Surender S. Samant. 2024. Hate Speech Detection in Social Media: Techniques, Recent Trends, and Future Challenges. *Wiley Interdisciplinary Reviews: Computational Statistics* 16, 2 (2024).
- [60] Pradeep K. Roy, Asis K. Tripathy, Tapan K. Das, and Xiao-Zhi Gao. 2020. A Framework for Hate Speech Detection Using Deep Convolutional Neural Network. *IEEE Access* 8 (2020), 204951–204962.
- [61] Haji M. Saleem, Kelly P. Dillon, Susan Benesch, and Derek Ruths. 2017. A Web of Hate: Tackling Hateful Speech in Online Social Spaces. arXiv:1709.10159 [cs.CL]
- [62] Vlad Sandulescu. 2020. Detecting Hateful Memes Using a Multimodal Deep Ensemble. arXiv:2012.13235 [cs.LG]
- [63] Shardul Suryawanshi, Bharathi R. Chakravarthi, Mihael Arcan, and Paul Buitelaar. 2020. Multimodal Meme Dataset (MultiOFF) for Identifying Offensive Content in Image and Text. In *Proceedings of the Second Workshop on Trolling, Aggression and Cyberbullying*. European Language Resources Association (ELRA), Marseille, France, 32–41. <https://aclanthology.org/2020.trac-1.6>
- [64] George-Alexandru Vlad, George-Eduard Zaharia, Dumitru-Clementin Cercel, and Mihai Dascalu. 2020. UPB@DANKMEMES: Italian Memes Analysis-Employing Visual Models and Graph Convolutional Networks for Meme Identification and Hate Speech Detection. *EVALITA Evaluation of NLP and Speech Tools for Italian* (2020), 288–293.
- [65] Hongbo Wang, Junyu Lu, Yan Han, Kai Ma, Liang Yang, and Hongfei Lin. 2025. Towards Patronizing and Condescending Language in Chinese Videos: A Multimodal Dataset and Detector. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1–5.
- [66] Han Wang, Rui Y. Tan, and Roy K. Lee. 2025. Cross-Modal Transfer from Memes to Videos: Addressing Data Scarcity in Hateful Video Detection. In *Proceedings of the ACM on Web Conference 2025* (Sydney NSW, Australia) (Www '25). Association for Computing Machinery, New York, NY, USA, 5255–5263. doi:10.1145/3696410.3714534
- [67] Han Wang, Rui Y. Tan, Usman Naseem, and Roy K. Lee. 2024. MultiHateClip: A Multilingual Benchmark Dataset for Hateful Video Detection on YouTube and Bilibili. In *Proceedings of the 32nd ACM International Conference on Multimedia* (Melbourne VIC, Australia) (MM '24). Association for Computing Machinery, New York, NY, USA, 7493–7502. doi:10.1145/3664647.3681521
- [68] Richard A. Wilson and Molly K. Land. 2020. Hate Speech on Social Media: Content Moderation in Context. *Connecticut Law Review* 52 (2020), 1029.
- [69] Ching S. Wu and Unnathi Bhandary. 2020. Detection of Hate Speech in Videos Using Machine Learning. In *2020 International Conference on Computational Science and Computational Intelligence (CSCI)*. IEEE, Las Vegas, NV, USA, 585–590. doi:10.1109/csci51800.2020.00104
- [70] Yang Wu, Pengwei Zhan, Yunjian Zhang, Liming Wang, and Zhen Xu. 2021. Multimodal Fusion with Co-Attention Networks for Fake News Detection. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online, 2560–2569. doi:10.18653/v1/2021.findings-acl.226
- [71] Haitao Xiong, Wei Jiao, and Yuanyuan Cai. 2024. TCE-DBF: Textual Context Enhanced Dynamic Bimodal Fusion for Hate Video Detection. *Data Technologies and Applications* (2024).
- [72] Fan Yang, Xiaochang Peng, Gargi Ghosh, Reshef Shilon, Hao Ma, Eider Moore, and Goran Predovic. 2019. Exploring Deep Multimodal Fusion of Text and Photo for Hate Speech Classification. Association for Computational Linguistics, Florence, Italy. doi:10.18653/v1/W19-3502
- [73] Weibo Zhang, Guihua Liu, Zhuohua Li, and Fuqing Zhu. 2020. Hateful Memes Detection via Complementary Visual and Linguistic Networks. arXiv:2012.04977 [cs.CV]

APPENDIX A. EARLY VS. LATE FUSION

Table 6 (model I), which corresponds to the results for CMA as a late fusion strategy; and Table 7, which contains the results for CMA as early fusion (model II), performing cross-modal attention directly on embeddings, followed by a feedforward layer. Both tables contain extensive runs on the different modality combinations and key-query pairs.

Table 6: Detailed multi-modal results for the late fusion, using CMA as a fusion strategy. The features are concatenated if more than one modality is used for a key and value. Key – K: key, Q: query, T: transcript, O: on-screen text, A: audio, V: video, M: macro average across both classes, H: hate class, F1: F1-score, ACC: unbiased accuracy, P: precision, R: recall.

Modality	K	Q	ACC	M-F1	F1(H)	P(H)	R(H)
TO	O	T	0.658	0.630	0.527	0.617	0.460
	T	O	0.829	0.825	0.800	0.776	0.825
TA	A	T	0.816	0.808	0.770	0.797	0.746
	T	A	0.816	0.811	0.781	0.769	0.794
TV	V	T	0.737	0.733	0.701	0.662	0.746
	T	V	0.829	0.825	0.797	0.785	0.810
OA	A	O	0.816	0.812	0.785	0.761	0.810
	O	A	0.829	0.820	0.780	0.836	0.730
OV	V	O	0.743	0.735	0.688	0.694	0.683
	O	V	0.763	0.752	0.700	0.737	0.667
AV	V	A	0.816	0.807	0.767	0.807	0.730
	A	V	0.822	0.812	0.769	0.833	0.714
TOA	OA	T	0.796	0.791	0.760	0.742	0.778
	TA	O	0.829	0.824	0.794	0.794	0.794
	TO	A	0.809	0.807	0.785	0.736	0.841
TOV	OV	T	0.770	0.759	0.711	0.741	0.683
	TV	O	0.855	0.852	0.828	0.815	0.841
	TO	V	0.829	0.824	0.797	0.784	0.809
TVA	VA	T	0.796	0.785	0.735	0.796	0.682
	TV	A	0.835	0.829	0.797	0.817	0.778
	TA	V	0.842	0.839	0.818	0.783	0.857
OAV	AV	O	0.796	0.789	0.752	0.758	0.746
	OV	A	0.809	0.803	0.768	0.774	0.762
	OA	V	0.809	0.805	0.775	0.757	0.794
TOAV	OAV	T	0.789	0.786	0.758	0.725	0.794
	TAV	O	0.882	0.877	0.852	0.881	0.825
	TOV	A	0.882	0.879	0.862	0.836	0.889
	TOA	V	0.842	0.837	0.810	0.810	0.810

Table 7: Detailed results for experiments using CMA as a fusion strategy on raw inputs. Key – K: key, Q: query, T: transcript, O: on-screen text, A: audio, V: video, M: macro average across both classes, H: hate class, F1: F1-score, ACC: unbiased accuracy, P: precision, R: recall.

Modality	K	Q	ACC	M-F1	F1(H)	P(H)	R(H)
TO	O	T	0.645	0.644	0.630	0.554	0.730
	T	O	0.829	0.825	0.800	0.776	0.825
TA	A	T	0.809	0.806	0.779	0.750	0.809
	T	A	0.829	0.825	0.797	0.785	0.810
TV	V	T	0.743	0.727	0.742	0.650	0.749
	T	V	0.836	0.832	0.809	0.779	0.841
OA	A	O	0.783	0.776	0.736	0.742	0.730
	O	A	0.632	0.631	0.622	0.541	0.730
OV	V	O	0.783	0.777	0.740	0.746	0.734
	O	V	0.651	0.651	0.634	0.561	0.730
AV	V	A	0.770	0.760	0.711	0.741	0.683
	A	V	0.796	0.789	0.752	0.758	0.746
TOA	OA	T	0.783	0.781	0.759	0.703	0.825
	TA	O	0.816	0.811	0.781	0.769	0.793
	TO	A	0.816	0.813	0.791	0.746	0.841
TOV	OV	T	0.776	0.770	0.730	0.730	0.730
	TV	O	0.842	0.839	0.818	0.783	0.857
	TO	V	0.829	0.825	0.800	0.776	0.825
TVA	VA	T	0.810	0.807	0.785	0.801	0.841
	TV	A	0.862	0.858	0.835	0.82813	0.841
	TA	V	0.816	0.812	0.785	0.761	0.810
OAV	AV	O	0.770	0.769	0.752	0.679	0.841
	OV	A	0.789	0.786	0.761	0.718	0.810
	OA	V	0.809	0.807	0.788	0.730	0.857
TOAV	OAV	T	0.822	0.820	0.797	0.757	0.841
	TAV	O	0.888	0.884	0.864	0.871	0.860
	TOV	A	0.855	0.851	0.825	0.851	0.825
	TOA	V	0.842	0.840	0.824	0.767	0.889

APPENDIX B. STOPWORDS

When removing stopwords from the transcript modality, we observe a small performance drop. The MM-HSD model with stopwords achieves an M-F1 of 0.874 and an F1-H of 0.853. After filtering out stopwords, M-F1 falls to 0.862, and F1-H to 0.841. This indicates that the model relies to some extent on certain high-frequency words from the transcript modality, such as negations, to identify HS. The full results are shown in Table 8.

Table 8: Effect of applying stopwords removal to transcript and OCR modalities in MM-HSD. Key – M: macro average across both classes, H: hate class, F1: F1-score, ACC: unbiased accuracy, P: precision, R: recall.

Model	ACC	M-F1	F1(H)	P(H)	R(H)
MM-HSD	.878 (.009)	.874 (.009)	.853 (.009)	.849 (.017)	.857 (.000)
MM-HSD (removing stopwords)	.866 (.006)	.862 (.006)	.841 (.006)	.826 (.011)	.857 (.000)

APPENDIX C. ANALYSIS OF MODALITIES INCLUDED IN CMA

Keeping the OCR as query, we experiment with removing modalities from the keys, while keeping all the modalities in the late fusion stage, with results presented in Table 9. We observe that MM-HSD achieves the best performance across all metrics, suggesting that different modalities add complementary information. The worst modality is text-only, possibly due to the fact that there is some informational overlap between transcript and OCR, as on-screen text is sometimes subtitles from the audio track. However, audio-only is the second-best model, suggesting that audio contains different information than OCR, such as volume and emotion.

Table 9: Performance of using excluding modalities from CMA. OCR is kept as the query modality in all cases. Key – M: macro average across both classes, H: hate class, F1: F1-score.

Modality	M-F1	F1(H)
Audio only	0.870	0.848
Video only	0.864	0.841
Transcript only	0.855	0.832
Audio + Video	0.866	0.845
Audio + Transcript	0.859	0.834
Video + Transcript	0.861	0.837
MM-HSD (A+V+T)	0.874	0.853