

Multimodal Learning using Optimal Transport for Sarcasm and Humor Detection

Shraman Pramanick*, Aniket Roy*, Vishal M. Patel

Johns Hopkins University

{spraman3, aroy28, vpatel136}@jhu.edu

Abstract

Multimodal learning is an emerging yet challenging research area. In this paper, we deal with multimodal sarcasm and humor detection from conversational videos and image-text pairs. Being a fleeting action, which is reflected across the modalities, sarcasm detection is challenging since large datasets are not available for this task in the literature. Therefore, we primarily focus on resource-constrained training, where the number of training samples is limited. To this end, we propose a novel multimodal learning system, **MuLOT** (Multimodal Learning using Optimal Transport), which utilizes self-attention to exploit intra-modal correspondence and optimal transport for cross-modal correspondence. Finally, the modalities are combined with multimodal attention fusion to capture the inter-dependencies across modalities. We test our approach for multimodal sarcasm and humor detection on three benchmark datasets - MUStARD [7] (video, audio, text), UR-FUNNY [20] (video, audio, text), MST [3] (image, text) and obtain 2.1%, 1.54% and 2.34% accuracy improvements over state-of-the-art.

1. Introduction

With an abundance of user-generated multimodal content, such as videos, multimodal learning has become an important area of research [59, 47, 2]. Unlike traditional unimodal learning on isolated modalities (such as vision, language, or acoustic), multimodal learning aims to aggregate complementary sources of information into a unified system. Understanding sarcasm from multimodal dialogues is a specialized form of sentiment analysis, where the speaker creatively experiments with words (language), gestures (vision), prosody (acoustic) to deliver incongruity across modalities. Humor is also a closely related quintessential sentiment often manifested using exaggeration or irony across modalities, such as a sudden twist of tone or a funny gesture. These two forms of subtle human sentiments are crucial to removing barriers

*denotes equal contribution

Time



Sheldon: It's just a privilege to watch your mind in work

Figure 1: Sample sarcastic utterance from the MUStARD dataset along with its transcript. Sheldon's comment (text) with a straight face (visual) and a clinical tone (acoustic) makes the instance sarcastic.

in conversations, building trust [55] and creating a positive impact on mental health [29]. However, existing deep neural systems often struggle to understand such fine-grained multimodal sentiments.

Motivation: The most common form of sarcasm and humor has traditionally been delivered using text. However, sarcasm in multimodal data often requires precise inter-modal cues to reveal the speaker's intentions. For instance, it can often be expressed using a combination of verbal and non-verbal cues, such as a change of tone, overemphasis in a word, a drawn-out syllable, or a straight-looking face. Consider the example shown in Figure 1 - the seemingly applauding utterance "It's just a privilege to watch your mind in work" becomes sarcastic when delivered with a straight face and clinical tone, overall expressing a negative connotation. Humans can naturally integrate all this simultaneous information subconsciously. However, building an algorithm that can potentially do the same requires appropriate representation of all these disparate sources of information and thus, gained immense research interest.

Challenges: Despite being regularly used in movies and sitcom shows, sarcastic conversational utterances are hard to collect, mainly because of the manual effort required in detecting and annotating raw videos with sarcasm labels. The only publicly available corpus, MUStARD, contains 690 video samples with an even number of sarcastic and non-sarcastic labels. As a consequence, off-the-shelf multimodal transformer-based systems [51, 37, 24, 47] with a

large number of parameters often tend to overfit this dataset. In this work, we propose MuLOT, which utilizes multi-head self-attention and optimal transport (OT) - [57] based domain adaptation to learn strong intra- and cross-modal dependencies using a small number of training samples. In order to embed asynchronous unimodal features in a common subspace, we use optimal transport kernel (OTK) Embedding [34]. Lastly, the self and cross-attended features are fused using an attention fusion module.

Contributions: In summary, our contributions are three-fold. (i) We propose a new multimodal learning method that uses multi-head self-attention and optimal transport (OT) - [57] based domain adaptation to learn strong intra- and cross-modal dependencies using a small number of training samples. (ii) We evaluate the proposed system on three different multimodal sarcasm and humor detection datasets and obtain 2.1 – 2.4% improvements over the previous state-of-the-art. (iii) We reduce the training corpus of two large multimodal datasets and demonstrate the superiority of our system over multimodal transformers in a resource-constrained training setup. On limited training samples, the proposed system outperforms state-of-the-art by an accuracy margin of 4.2 – 5.8%.

2. Related Work

The proliferation of multimedia data on the Internet has resulted in multimodal systems gaining an increasing interest in recent years. One of the fundamental challenges in a multimodal framework is to fuse different unimodal features, which are asynchronous and have varying dimensions. Poria et al. [42] presented an outline of various fusion techniques and potential performance improvements with multiple modalities.

Multimodal Fusion: Similar to human behavior, multimodal frameworks aim to integrate and correlate simultaneous multiple resources, such as acoustic, visual, and textual information [7, 21, 2, 45]. There are mainly three types of fusion strategies – early, late, and hybrid. Early fusion methods [43, 64] integrates different sources of data into a single feature vector and uses a single classifier. On the other hand, late fusion [4, 5] aggregates the outputs of different classifiers trained on different modalities. However, none of these techniques consider the interdependence among the different modalities and hence perform poorly on real applications. In contrast, hybrid fusion [65, 1, 58, 41, 44, 46] jointly learns from various modality-specific sources by employing an intermediate shared representation layer and has been most prominent in the literature. Recently following the success of transformers [56, 23, 40, 14, 63, 31], several works have extended it to multimodal applications by capturing both unimodal and cross-modal interactions via fine-tuning [54, 33, 51, 37, 24, 25, 47]. For example, Rahman et al. [47] integrate acoustic and visual information in

pre-trained transformers like BERT [14] and XLNet [63]. However, all the multimodal transformers require a large number of training samples for good performance.

Multimodal Sarcasm & Humor Detection: Humor and sarcasm are two closely related sentiments - while humor is often expressed using exaggeration and irony, sarcasm mostly generates from incongruity. Unimodal humor and sarcasm detection are well-studied in literature [53, 60, 48, 62, 8, 30, 10, 28]. However, the multimodal counterpart is more subtle as the irony or incongruity is often present in different modalities. Hasan et al. [20] introduced the first large-scale multimodal dataset (UR-FUNNY) for humor detection. Castro et al. [7] collected a multimodal conversational dataset (MUSARD) for sarcasm detection from popular sitcom TV shows. Cai et al. [3] developed another dataset to detect multimodal sarcasm in Twitter posts. Recently, Han et al. [19] proposed an end-to-end network that performs fusion (relevance increment) and separation (difference increment) on pairwise modality representations for humor detection. Another closely related work, MISA [21] aggregates modality-invariant and modality-specific representations and has been applied to predict humor in the UR-FUNNY dataset. However, multimodal sarcastic video samples are hard to collect, and the existing neural systems struggle to perform well on MUSARD. In this paper, we utilize self-attention to capture intra-modal correspondence and optimal transport for cross-modal correspondence in a low-resource training setup.

3. MuLOT: Proposed System

This section describes our proposed system, MuLOT for sarcasm and humor detection in videos and images by leveraging multimodal signals. Each video consists of multiple utterances¹, where each utterance - a smaller video by itself, is considered to contain sarcastic/humorous dialogues for positive samples. These utterances are a multimodal source of data - we preprocess the input utterances to extract fine-grained unimodal features from three modalities - visual (v), language (l), and acoustic (a). Additionally, every utterance is accompanied by a context - which is also a short video helping to discern the background of the utterance. We concatenate the context with the utterance so that the utterance can be modeled in light of the context and the resulting video is treated as the input to the system. Consequently, our system is generalizable to images. Each input image is associated with a caption and a piece of text on it², which we consider as visual (v), caption text (t_c) and OCR text (t_{ocr}) modalities. Features corresponding to these modalities are represented as $U_m \in \mathbb{R}^{L_m \times d_m}$, where L_m denotes the sequence length for modality m and d_m is the

¹An utterance is a unit of speech bounded by breaths or pauses [36]

²We employ Google OCR Vision API to extract this text

respective feature dimension. The details of these features are discussed in Section 3.1. Given these unimodal feature sequences $U_{m \in \{v, l/t_c, a/t_{ocr}\}}$, the task is to predict the affective orientation of the input in a binary classification setup - sarcastic/humorous or non-sarcastic/non-humorous.

As shown in Figure 2, MuLOT consists of two modules, the self-attention module and the cross-attention module, shown in red and blue dashed blocks, respectively. First, the modality-specific fine-grained features are fed into the self-attention module which aims to learn intra-modality correspondence by emphasizing the most relevant tokens in the entire feature sequence. In parallel, the cross-attention module maps pairwise unimodal modalities into a common feature space, and thus learns inter-modality interaction. By taking both intra-modality and inter-modality relationships into consideration, the system gains the ability to synchronize across different modalities over various time and learns the correlation among them [59]. Finally, the self-attended unimodal features and cross-attended shared features are fused by using an attention fusion mechanism. It is important to observe that MuLOT achieves the same objective as multimodal transformers [54, 33, 51, 37, 24, 25, 47], but has significantly fewer trainable parameters, and thus works better in low-resource training setup.

3.1. Unimodal Feature Extraction

In the proposed multimodal learning setup, we leverage three modalities, e.g., video, audio and text. Initially we extract fine-grained unimodal features from these modalities as described below.

Visual Features: The video features are extracted using a pre-trained action recognition model, *I3D* [6] trained on the Kinetics dataset [26] to recognize 400 different human actions. All the frames, computed at a rate of 5 FPS, are first preprocessed by resizing, center-cropping, and normalization to have a resolution of 112×112 . For every 16 non-overlapping frames in a video, *I3D* extracts a 2048 dimensional feature vector. Therefore, the final unimodal feature dimension for every video is $L_v \times 2048$, where L_v is the number of sets 16 non-overlapping frames. For images, we employ ResNet-101 [22] pre-trained on ImageNet [13]. Specifically, we extract $7 \times 7 \times 2048$ feature maps from the last pooling layer (pool5) of ResNet-101 and reshape it into a dimension of 49×2048 . Hence, for the images, $L_v = 49$.

Language Features: Traditionally, language modality features have been GloVe [39] or Word2Vec [35] embeddings for each token. However, following recent works [7, 52, 21, 37], we utilize the pre-trained BERT [14] as the feature extractor for textual utterances. In particular, we input the raw text to a pre-trained uncased BERT-base model to get a 768-dimensional dynamic feature representation for every token in the utterance. Hence, the resulting language feature dimension is $L_l \times 768$, where L_l is the number of

tokens in the utterance.

Acoustic Features: The acoustic modality is expected to contribute information related to pitch, intonation, and other tonal-specific details of the speaker [53]. To achieve this, we obtain low-level features from the audio stream for each video. Following state-of-the-art, we use COVAREP [12] to extract the audio features. These low-level features include but not limited to 12 Mel-frequency cepstral coefficients, Voiced/Unvoiced segmenting features (VUV) [15], glottal source parameters [16]³. The resulting acoustic feature dimension is $L_a \times 81$, where L_a depends on the length of the video. Note that we use the OT kernel module in our system to align all three temporal sequences within an utterance to be of equal length, i.e. $L_v = L_l = L_a$.

3.2. Intra-Modality Attention

Now we introduce the self-attention module used to model the intra-modal correspondence among three different modalities. Typically, an attention module can be described as a tuple of key, query and value, where the output is a weighted sum of the values and the weight matrix is determined by query and its corresponding key. In case of self-attention, the key, query and value are equal. Following the philosophy of [56], the recent trend is to apply transformers for modeling self-attention. However, a transformer module consists of multiple sub-layers and is often difficult to train from scratch. In our low-resource training setup, we simplify the structure of a transformer module by utilizing the multi-head attention layers.

We compute the multi-head self-attention as follows. Given the unimodal feature matrices $U_m \in \mathbb{R}^{L_m \times d_m}$, where $m \in \{v, l/t_c, a/t_{ocr}\}$, we aim to extract k distinct set of relevant features corresponding to the input representation, where k is a hyper-parameter denoting the number of attention heads. The multi-head attention mechanism takes U_m as an input, and outputs an attention weight matrix $W_m \in \mathbb{R}^{k \times L_m}$, respective to the modality-specific representation as follows

$$W_m = \text{softmax}(W_{h2} \tanh(W_{h1} U_m^T)), \quad (1)$$

where, $W_{h1} \in \mathbb{R}^{r \times d_m}$ and $W_{h2} \in \mathbb{R}^{k \times r}$ are parameter matrices to be learned during training. The $\text{softmax}(\cdot)$ is performed along the second dimension of its input, and r is a hyper-parameter we can set arbitrarily. The resulting modality-specific embedding matrices U_{mm} are computed using their respective attention weights and unattended features as follows

$$U_{mm} = W_m U_m, \quad (2)$$

where, $U_{mm} \in \mathbb{R}^{k \times d_m}$ are the resulting self-attended features for modality m .

³For the full set of features, please refer to [66, 20]

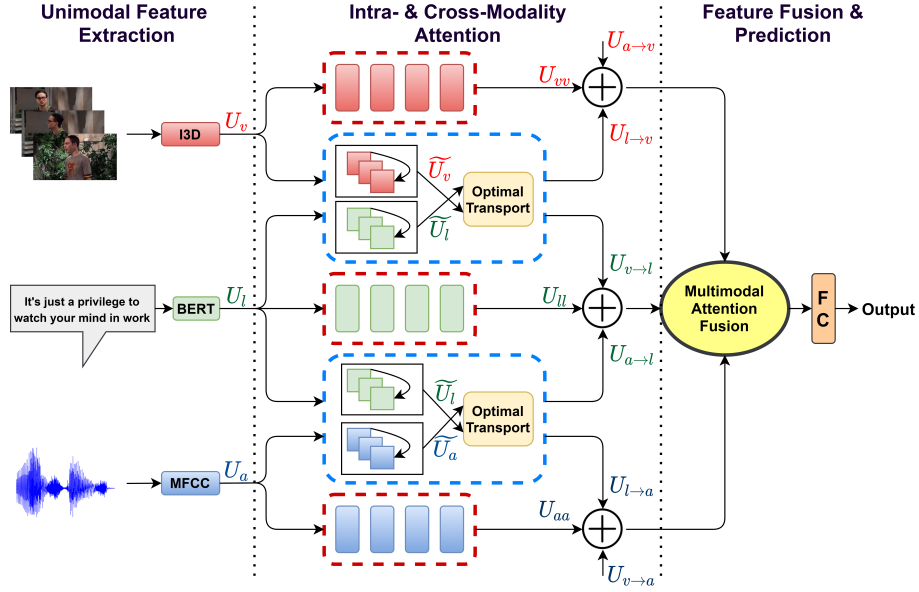


Figure 2: The architecture of the proposed model, MuLOT . The features from each modality are passed through multi-head self-attention module (red dashed) and optimal transport-based cross-attention module (blue dashed). Finally, the attended features are fused using Multimodal Attention Fusion (MAF).

3.3. Cross-Modality Attention

Although the self-attention module effectively exploits the intra-modality correspondence, the inter-modality relationship, i.e. the interdependence across three different modalities is not deployed. Motivated by the success of Transformers [56, 23, 40, 14, 63, 31], several recent works have extended it to learn multimodal association by adding additional vision and speech modules to the Transformer framework [54, 33, 51, 37, 24, 25, 47]. Although multimodal transformers rely on a language-pretrained BERT [14], which is fixed throughout, the vision and acoustic components have to be trained from scratch, which often requires hundreds of thousands of training samples. To this end, instead of using a transformer-based cross-attention module, we exploit a recently proposed technique viz., optimal transport kernel (OTK) [34] for modelling the inter-modality association. OTK combines the idea of kernel methods and optimal transport to fuse the unimodal features with varying dimension and dependencies into a same reference frame. The complete cross-attention module consists of two steps – unimodal feature transform and domain adaptation using optimal transport.

Unimodal Feature Transform: Since the visual, language and acoustic features have different dimensions with varying sizes and dependencies, mapping those into a same reference frame is necessary yet challenging. In the OTK approach, the unimodal feature vectors are embedded into a reproducing kernel Hilbert space (RKHS) and then a weighted pooling operation is performed based on attention with weights provided by the transport plan between the set and a trainable

reference. The primary motivation for using kernels is to provide non-linear transformation of the input features. Afterward, optimal transport aligns the features on a trainable reference frame.

Suppose we want to compute an optimal transport plan from x to x' . Let a and b be defined as $a = \sum_i a_i \delta_i$ and $b = \sum_i b_i \delta_i$ and C be the pairwise costs for aligning the elements x to x' . Then the entropic regularized Kantorovich relaxation of OT from x to x' is [57]:

$$\min_{P \in U(a,b)} \sum_{ij} C_{ij} P_{ij} - \epsilon H(P), \quad (3)$$

where $H(P) = -\sum_{ij} P_{ij} (\log(P_{ij} - 1))$, is the entropic regularization with parameter ϵ , which controls the sparsity of P and U is the subspace of admissible coupling between a and b defined as follows

$$U(a,b) = \{P \in R_+ : P1_n = a \text{ and } P^T 1_{n'} = b\}. \quad (4)$$

Now, suppose $x = (x_1, x_2, \dots, x_n)$ be input feature vector and $z = (z_1, z_2, \dots, z_k)$ be the reference set with k elements. Let κ be a positive definite kernel and ϕ its corresponding kernel embedding. The κ matrix contains the values of $\kappa(x_i, z_i)$ before alignment and the transportation plan between x and z is denoted as $P(x, z)$, which is defined by the unique solution of Eq. (3) when choosing cost $C = -\kappa$. The embedding is defined as: $\Phi_z(x) = \sqrt{p} P(x, z)^T \phi(x)$, where $\phi(x) = [\phi(x_1), \dots, \phi(x_n)]^T$. $\Phi_z(x)$ can also be represented as a positive semidefinite kernel as follows [34]

$K_z(x, x') = \sum P_z(x, x')_{ij} \kappa(x_i, x_j) = \langle \Phi_z(x), \Phi_z(x') \rangle$ where, K_z is the OTK. Using OTK, we transform $L_m \times d_m$ dimensional unimodal features to a uniform reference frame with length L_{uni} and feature dimension d_{uni} , where

Dataset	#Samples	Train			Dev			Test		
		#Sar	#Non-Sar	Total	#Sar	#Non-Sar	Total	#Sar	#Non-Sar	Total
MUStARD	690	275	275	550	35	35	70	35	35	70
MST	24635	8642	11174	19816	959	1451	2410	959	1450	2409
Tiny MST	7819	2000	2000	4000	959	1451	2410	959	1450	2409
		#Hum	#Non-Hum	Total	#Hum	#Non-Hum	Total	#Hum	#Non-Hum	Total
UR-FUNNY	16514	5306	5292	10598	1313	1313	2626	1638	1652	3290
Tiny UR-F	8916	1500	1500	3000	1313	1313	2626	1638	1652	3290

Table 1: Statistics about the five datasets used in our experiments. We have reduced the training set of MST and UR-FUNNY to generate Tiny MST and Tiny UR-F, respectively. However, the corresponding dev and test sets remain unchanged.

L_{uni} and d_{uni} are chosen by experiments. The transformed features are represented as $\tilde{U}_m \in \mathbb{R}^{L_{uni} \times d_{uni}}$, where $m \in \{v, l/t_c, a/t_{ocr}\}$.

Domain Adaptation using Optimal Transport: After transforming the varying dimensional unimodal features into a uniform reference frame, we finally embed pairwise unimodal features into a common subspace, using the OT-based domain adaptation. To implement OT as part of a larger learning model, we use the recently released Python optimal transport library [17] in our experiments.

We transport among each pair of modalities, which can be interpreted as domain adaptation across two modalities and the transported features are then concatenated to obtain the final adapted features. $U_{n \rightarrow m}$ denotes the transported features from modality n to modality m , i.e.

$$U_{n \rightarrow m} = OT(\tilde{U}_n \rightarrow \tilde{U}_m). \quad (5)$$

We then concatenate transported features to the self-attended features U_{mm} to obtain the final adopted features as follows

$$U_{shared}^m = U_{mm} \oplus U_{n \rightarrow m} \oplus U_{p \rightarrow m}, \quad (6)$$

where, $m, n, p \in \{v, l/t_c, a/t_{ocr}\}$, $m \neq n \neq p$, and \oplus denotes simple concatenation.

3.4. Multimodal Attention Fusion (MAF)

MAF utilizes an attention-based mechanism to fuse three set of features, U_{shared}^m , corresponding to the three modalities. For some samples, the visual modality is relevant, while for others, the language or acoustic modality plays more crucial role. Hence, the MAF module aims to model the relative important of different modalities. Motivated by [18, 44], we design our MAF module with two major parts – modality attention generation and weighted feature concatenation. In the first part, a sequence of dense layers followed by a softmax layer is used to generate the attention scores $[w_v, w_l, w_a]$ for the three modalities. In the second part, the adopted features are weighted using their respective attention scores and concatenated together as follows

$$U_{final}^v = (1 + w_v)U_{shared}^v \quad (7)$$

$$U_{final}^l = (1 + w_l)U_{shared}^l \quad (8)$$

$$U_{final}^a = (1 + w_a)U_{shared}^a \quad (9)$$

$$U_{final}^{v,l,a} = W_U \otimes [U_{final}^v, U_{final}^l, U_{final}^a] \quad (10)$$

We also use residual connections for better gradient flow. The final multimodal representation, $U_{final}^{v,l,a}$, is fed into a series of fully-connected layers for the final binary classification.

4. Experiments

In this section, we present details of the datasets used, unimodal feature extraction steps, baselines, and training methodologies.

4.1. Datasets

We evaluate MuLOToN three different benchmark datasets for multimodal sarcasm and humor detection. Additionally, we access the effectiveness of MuLOToN in a low-resource setup by shrinking the large training sets of two datasets.

MUStARD: Multimodal Sarcasm Detection Dataset (MUS-tARD) [7] is the only available resource to enable sarcasm detection in conversational videos. This dataset is curated from popular TV shows like Friends, The Big Bang Theory and consists of audiovisual utterances annotated with sarcasm labels. Each target utterance in this dataset is associated with historical dialogues as context, which is key to understanding the backdrop of sarcastic remarks. The primary challenge in using MUStARD is its small size, which limits the performance of heavy transformer-based systems on this corpus.

MST: Multimodal Sarcasm in Twitter Posts (MST) [3] consists of sarcastic and non-sarcastic image-text pairs collected from Twitter. This dataset is primarily bimodal, as there is no acoustic modality. However, we notice that a significant amount of samples contains textual information on the image. We employ Google OCR Vision API⁴ to extract this text and treat this as the third modality. We further reduce the training set of MST and name the resulting corpus Tiny MST.

UR-FUNNY: For multimodal humor detection, we use the UR-FUNNY dataset [20] which is collected from TED talk videos and therefore has three modalities. Similar to MUS-tARD, this dataset consists of context preceding the target punchline. We further create a tiny variant of UR-FUNNY

⁴<https://cloud.google.com/vision/docs/ocr>

by truncating its training set. The split statistics of all these datasets are presented in Table 1.

4.2. Baselines

For unified comparison across videos and images, we use the following baselines. Furthermore, we remove various modalities and modules at a time from the proposed MuLOT system to observe the effect in performance. Baselines on MUSTARD and UR-FUNNY include - **Support Vector Machines (SVM)** [11], **DFF-ATMF** [9], **CIM-MTL** [1], **Tensor Fusion Network (TFN)** [64], **Contextual Memory Fusion Network (CMFN)** [20], **MISA** [21], **MAG-XLNet** [47]. MAG-XLNet introduced Multimodal Adaption Gate (MAG) to fuse acoustic and visual information in pre-trained language transformers, and is the SOTA on both MUSTARD and UR-FUNNY. Baselines on MST include- **Concat BERT** (concatenates the features extracted by pre-trained unimodal ResNet-152 [22] and Text BERT [14] and uses a simple perceptron as the classifier), **Supervised Multimodal Bi-transformer (MMBT)** [27], **Vision and Language BERT (ViLBERT)** [33], **Hierarchical Fusion Model (HFM)** [3], **D&R Net** [61], **MsdBERT** [37]. MsdBERT, which exploits a co-attention network to exploit intra- and inter-modality incongruity between text and image, is the SOTA method on the MST dataset. More details on the baselines are provided in supplementary document.

4.3. Training

We train MuLOT using Pytorch framework [38] on Nvidia-RTX 2080Ti GPUs, with 24 GB dedicated memory in each GPU. As described in Section 3.1, we use pre-trained unimodal feature extraction models, and fine-tune them during training. All other weights are randomly initialized with a zero-mean Gaussian distribution with standard deviation 0.02.

Although MUSTARD and UR-FUNNY are balanced datasets, Table 1 shows label imbalance for the MST dataset. Therefore, when training on MST dataset, we assign larger weight ($[w_{sar}, w_{non-sar}] = [1.2, 1]$) for minority class to minimize label imbalance. We train our model using the Adam [32] optimizer and the binary cross-entropy loss as the objective function. The initial learning rate is 0.005 and the network is trained for 300 epoches. The detailed hyper-parameters used for the training is provided in the supplementary material.

5. Results, Discussion and Analysis

In this section, we compare the performance of MuLOT system with different multimodal baselines, conduct a detailed ablation study to demonstrate the importance of different modalities and modules in our system. Furthermore, we visualize the interpretability of MuLOT using Grad-CAM [49]. Since the test sets of MUSTARD and UR-FUNNY are

Algorithm	Context	Target	MUSTARD	UR-FUNNY	Tiny UR-F
			Acc \uparrow	Acc \uparrow	Acc \uparrow
SVM	\times	\checkmark	73.55	-	-
DFF-ATMF	\times	\checkmark	64.45	62.55	56.35
CIM-MTL	\times	\checkmark	67.14	63.20	56.71
TFN	\times	\checkmark	68.57	64.71	57.23
CMFN (GloVe)	\times	\checkmark	67.14	64.47	57.10
CMFN (GloVe)	\checkmark	\checkmark	70.00	65.23	59.25
MISA (BERT)	\times	\checkmark	66.18	70.61	62.66
BBFN	\checkmark	\checkmark	71.42	71.68	63.20
MAG-XLNet	\checkmark	\checkmark	74.72	72.43	67.22
MuLOT	\times	\checkmark	74.52	73.22	70.74
MuLOT	\checkmark	\checkmark	76.82\dagger	73.97\dagger	71.46\dagger
$\Delta_{\text{MuLOT}-\text{baseline}}$			$\uparrow 2.10$	$\uparrow 1.54$	$\uparrow 4.24$

Table 2: Performances of multimodal models on the MUSTARD, UR-FUNNY and Tiny UR-F datasets. Since the test sets are balanced (c.f. Table 1), only binary classification accuracy is reported. \dagger indicates the results have experienced paired t-test with $p < 0.01$ and demonstrate significant improvement over MAG-XLNet, the best baseline model.

Algorithm	OCR	MST		Tiny MST	
		Acc \uparrow	F1 \uparrow	Acc \uparrow	F1 \uparrow
Concat BERT	\times	81.08	79.56	76.21	73.48
HFM	\times	83.44	80.18	77.80	74.07
D&R Net	\times	84.02	80.60	79.43	76.72
MMBT	\times	83.46	80.74	79.48	76.09
MMBT	\checkmark	84.87	82.66	80.57	77.20
ViLBERT	\times	84.21	82.49	79.42	75.95
ViLBERT	\checkmark	86.90	84.22	80.68	77.24
MsdBERT	\times	86.05	82.92	80.14	77.53
MsdBERT	\checkmark	88.75	86.18	82.30	79.90
MuLOT	\times	87.41	86.33	84.46	82.62
MuLOT	\checkmark	90.82\dagger	88.52\dagger	88.04\dagger	85.93\dagger
$\Delta_{\text{MuLOT}-\text{baseline}}$		$\uparrow 2.07$	$\uparrow 2.34$	$\uparrow 5.74$	$\uparrow 6.03$

Table 3: Performance of multimodal models on the MST and Tiny MST datasets. Since the test set is imbalanced (c.f. Table 1), both binary classification accuracy and macro F1 scores are reported. \dagger indicates the results have experienced paired t-test with $p < 0.01$ and demonstrate significant improvement over MsdBERT, the state-of-the-art model.

balanced, we use binary accuracy as the evaluation metric. However, since the test set of MST is imbalanced, we also consider F1 scores when evaluating on this dataset.

5.1. Comparison with Baselines

Multimodal Sarcasm Detection: Table 2 presents the comparative classification performances on the MUSTARD dataset. Since this corpus is relatively small, complex neural models with a large number of parameters often overfit when trained from scratch. The authors of the MUSTARD paper reported the best accuracy of 71.60% when using a simple SVM as classifier [7]. However, they used framewise ResNet features as the visual modality, which does not consider the temporal dynamics of the videos. We retrain an SVM using I3D features and produce 73.55% accuracy to indicate the effectiveness of I3D features for videos. Transformer-based

Modality	Algorithm	MUStARD	UR-FUNNY	Tiny UR-F	Algorithm	MST		Tiny MST	
		Acc \uparrow	Acc \uparrow	Acc \uparrow		Acc \uparrow	F1 \uparrow	Acc \uparrow	F1 \uparrow
Trimodal	MuLOT	78.57	73.97	71.46	MuLOT	90.82	88.52	88.04	85.93
Unimodal	visual only	73.30	60.72	58.80	visual only	82.65	81.22	78.56	77.70
	language only	73.54	69.58	67.32	caption only	83.40	82.14	80.06	78.85
	acoustic only	64.00	64.35	55.44	OCR only	78.64	77.39	76.22	75.31
Bimodal	visual + language	77.18	70.40	69.40	visual + caption	87.35	85.93	83.94	82.45
	visual + acoustic	75.54	69.23	69.82	visual + OCR	85.66	84.37	82.30	81.38
	language + acoustic	75.72	72.10	69.12	caption + OCR	85.10	83.79	81.87	81.00
Trimodal	MuLOT w/o self-att ⁿ	71.60	64.46	63.22	MuLOT w/o self-att ⁿ	86.24	84.80	83.69	82.84
	MuLOT w/o cross-att ⁿ	63.88	60.08	57.15	MuLOT w/o cross-att ⁿ	82.32	80.20	80.07	79.28
	MuLOT w/o MAF	75.23	71.22	68.84	MuLOT w/o MAF	87.94	86.73	85.55	84.38

Table 4: Ablation Study. Role of different modalities and various modules in our proposed MuLOT system on five datasets.

inter-utterance contextual baselines, DFF-ATMF and CIM-MTL, perform poorly on this dataset. We also observe the role of historical context dialogues to model sarcasm. The same CMFN baseline yields 2.86% better accuracy after using context. Nevertheless, baselines like MISA and TFN do not improve from context. As described by the respective authors, these models suffer in encoding the long sequence and thus, lose crucial information during fusion.

MAG-XLNet, which is the current state-of-the-art model in CMU-MOSI and CMU-MOSEI datasets for multimodal sentiment analysis, also performs well on MUStARD. Unlike other multimodal transformers, MAG-XLNet makes no change to the original structure of XLNet, but rather comes as an attachment to utilize multimodal information and hence proves relatively more effectual in a low resource setup. Our proposed system, MuLOT, improves on MAG-XLNet by 2.10%. It is important to note that, MuLOT has only 21M trainable parameters, which is almost 8 times smaller than MAG-XLNet. Domain adaptation using OT and low-number of parameters helps MuLOT in achieving the best results.

Table 3 shows the performance of baselines and MuLOT on the MST dataset. Since a significant amount of the image samples of this dataset have text on it, the state-of-the-art system, MsdBERT utilizes OCR extracted texts to improve its F1 score by 3.26%. MMBT and ViLBERT also show a very similar trend. Since the MST dataset has a large training corpus, sophisticated transformer-based systems significantly outperform simple models like Concat BERT and HFM. Our proposed system, MuLOT beats the SOTA on MST by an F1 score of 2.34%, demonstrating its effectiveness on large corpus as well. Furthermore, in order to evaluate MuLOT on a low amount of training data, we randomly reduce the training set of MST to have only 2000 samples in each class and call it Tiny MST. When trained on Tiny MST, the performance of MMBT, ViLBERT, and MsdBERT drops drastically. However, MuLOT is able to maintain an F1 score of 85.93% and beats the best baseline by a substantial margin of 6.03%. The low number of trainable parameters and optimal transport-based domain adaptation setup in MuLOT again proves helpful when trained with limited data.

Multimodal Humor Detection: For humor detection, we present the classification performances of baselines and

MuLOT in Table 2. Similar to MUStARD, MAG-XLNet produces the best baseline performance on the UR-FUNNY dataset. The inclusion of context also helps many baselines to boost their performance by a few points. MuLOT beats MAG-XLNet marginally on UR-FUNNY. However, when the training corpus size is reduced, MAG-XLNet drops its performance by 5.21% compared to MuLOT’s 2.51%. Thus, on all three low-resource setups, MuLOT consistently outperforms all the baselines by a large margin, demonstrating that multimodal transformers are not good enough in the absence of ample training samples.

5.2. Ablation Study

Role of Modalities: In Table 4, we retrain our system after removing one modality at a time to observe the effect in performance. First, we see that the trimodal system produces the best performance compared to all bimodal/unimodal systems, indicating that each modality contains complementary information. In the case of MUStARD, we observe similar drop in performance by removing either visual or language modalities, while for UR-FUNNY, the removal of language hurts the most. Since the cameras move a lot and seldom focuses on the faces in the TED videos of UR-FUNNY, the videos in this dataset carry a lot of noise.

Next, we retrain the MuLOT system using only one modality at a time to observe how much modality-specific information our unimodal encoders can capture. Following previous results, for MUStARD, both visual and language modality performs equally well, and for UR-FUNNY, the language modality proves to be the best. Audio modality alone can not yield good results because the audio features only capture the tonal-specific details of the speaker, such as pitch and intonation, which only makes sense when the transcript is present. For the MST dataset, the image and caption together perform the best. The OCR text is only present in around 57.2% of the samples and thus contributes less than other modalities. Figure 3 shows the complementarity of different modalities in each dataset. We extract the predictions from each unimodal encoder and calculate the Pearson correlation among them. The low correlation values, shown in lighter shades, indicate that each modality covers different aspects of information. The trimodal MuLOT system integrates all




Figure 3: Pearson correlation calculated among the prediction outputs (dev & test combined) of unimodal models.

these modalities in a unified framework and proves its efficacy.

Role of Intra- & Cross-modal Attention: Incorporating intra- and cross-modal attention significantly improves the performance, as shown in Table 4. Since the incongruity, irony, and exaggeration of sarcasm and humor can generate from different modalities, cross-modal learning becomes the key, as illustrated in Figure 1.

Role of MAF: Since the importance of three modalities can vary for different samples, we utilize an attention-based mechanism (MAF) to fuse three different sets of features. As shown in Table 4, we observe that compared to a simple concatenation of the modalities, MAF improves the performance by 2 – 4% across the five datasets. We also find that a 3-layer MAF network is sufficient to provide the best performance.

5.3. Interpretability of MuLOT

In this section, we comprehend the interpretability of MuLOT by generating the visual explanations over the videos and images by using Grad-CAM [49] and visualize the attention distribution over text. Figure 4 shows three different frames of a sarcastic utterance (video id: 1_60) from the MUSTARD dataset. MuLOT focuses on the facial expression (straight face) of the speaker for all frames. In the language modality, the system focuses on words like "privilege" and thus, can detect the incongruity across the modalities. Figure 5 shows a sarcastic tweet from the MST dataset. In this sample, OCR successfully detects the text present on the image and the system perceives the irony between the OCR extracted text and the caption. Moreover, the class-discriminative regions on the image are properly identified by MuLOT , as shown by Grad-CAM.

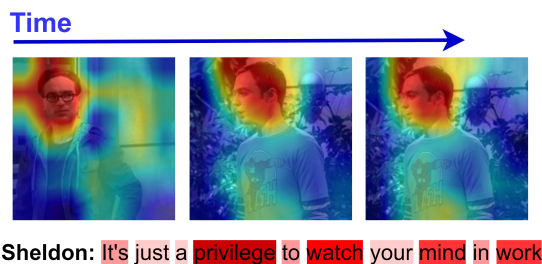


Figure 4: Visual explanations and textual attention map for a sarcastic utterance from the MUSTARD dataset.



Figure 5: Visual explanations and textual attention maps for a sarcastic tweet from the MST dataset.

6. Conclusion

In this paper, we deal with the task of detecting multi-modal sarcasm and humor from conversational videos and image-text pairs. Capturing the intra- and inter-modal dynamics for these actions, which are highly dependent on the synchronization and aggregation of different modalities, is quite challenging with limited training data. To this end, we propose MuLOT , which captures intra-modal dynamics using multi-head self-attention and cross-modal dynamics using optimal transport. Finally, multimodal attention fusion across the modalities has been performed, which further improves the performance. Experimental results on three benchmark datasets- MUSTARD (video, audio, text), UR-FUNNY (video, audio, text), MST (image, text) shows 2.1%, 1.54% and 2.34% performance improvements compared to the state-of-the-art using our proposed method.

Acknowledgement

The authors would like to thank Prof. Rama Chellappa for his helpful suggestions. Aniket Roy was supported by an ONR MURI grant N00014-20-1-2787 and Vishal M. Patel was supported by the NSF grant 1910141.

References

- [1] Md Shad Akhtar, Dushyant Chauhan, Deepanway Ghosal, Soujanya Poria, Asif Ekbal, and Pushpak Bhattacharyya. Multi-task learning for multi-modal emotion recognition and sentiment analysis. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 370–379, 2019.

- [2] Yash Kumar Atri, Shraman Pramanick, Vikram Goyal, and Tanmoy Chakraborty. See, hear, read: Leveraging multi-modality with guided attention for abstractive text summarization. *Knowledge-Based Systems*, page 107152, 2021.
- [3] Yitao Cai, Huiyu Cai, and Xiaojun Wan. Multi-modal sarcasm detection in Twitter with hierarchical fusion model. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2506–2515, Florence, Italy, July 2019. Association for Computational Linguistics.
- [4] Erik Cambria, Dipankar Das, Sivaji Bandyopadhyay, and Antonio Feraco. Affective computing and sentiment analysis. In *A practical guide to sentiment analysis*, pages 1–10. Springer, 2017.
- [5] Donglin Cao, Rongrong Ji, Dazhen Lin, and Shaozi Li. A cross-media public sentiment analysis system for microblog. *Multimedia Systems*, 22(4):479–486, 2016.
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [7] Santiago Castro, Devamanyu Hazarika, Verónica Pérez-Rosas, Roger Zimmermann, Rada Mihalcea, and Soujanya Poria. Towards multimodal sarcasm detection (an _Obviously_ perfect paper). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4619–4629, Florence, Italy, July 2019. Association for Computational Linguistics.
- [8] Arjun Chandrasekaran, Ashwin K Vijayakumar, Stanislaw Antol, Mohit Bansal, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. We are humor beings: Understanding and predicting visual humor. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4603–4612, 2016.
- [9] Feiyang Chen, Ziqian Luo, Yanyan Xu, and Dengfeng Ke. Complementary fusion of multi-features and multi-modalities in sentiment analysis. *arXiv preprint arXiv:1904.08138*, 2019.
- [10] Peng-Yu Chen and Von-Wun Soo. Humor recognition using deep learning. In *Proceedings of the 2018 conference of the north american chapter of the association for computational linguistics: Human language technologies, volume 2 (short papers)*, pages 113–117, 2018.
- [11] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.
- [12] Gilles Degottex, John Kane, Thomas Drugman, Tuomo Raitio, and Stefan Scherer. Covarep—a collaborative voice analysis repository for speech technologies. In *2014 IEEE international conference on acoustics, speech and signal processing (icassp)*, pages 960–964. IEEE, 2014.
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- [15] Thomas Drugman and Abeer Alwan. Joint robust voicing detection and pitch estimation based on residual harmonics. In *INTERSPEECH 2011, Twelfth Annual Conference of the International Speech Communication Association*, 2011.
- [16] Thomas Drugman, Mark Thomas, Jon Gudnason, Patrick Naylor, and Thierry Dutoit. Detection of glottal closure instants from speech signals: A quantitative review. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(3):994–1006, 2011.
- [17] Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, Léo Gautheron, Nathalie T.H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine Rolet, Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong, and Titouan Vayer. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021.
- [18] Yue Gu, Kangning Yang, Shiyu Fu, Shuhong Chen, Xinyu Li, and Ivan Marsic. Hybrid attention based multimodal network for spoken language classification. In *ACL*, volume 2018, page 2379, 2018.
- [19] Wei Han, Hui Chen, Alexander Gelbukh, Amir Zadeh, Louis-philippe Morency, and Soujanya Poria. Bi-bimodal modality fusion for correlation-controlled multimodal sentiment analysis. *23rd ACM International Conference on Multimodal Interaction (ICMI)*, 2021.
- [20] Md Kamrul Hasan, Wasifur Rahman, AmirAli Bagher Zadeh, Jianyuan Zhong, Md Iftexhar Tanveer, Louis-Philippe Morency, and Mohammed (Ehsan) Hoque. UR-FUNNY: A multimodal language dataset for understanding humor. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2046–2056, Hong Kong, China, Nov. 2019. Association for Computational Linguistics.
- [21] Devamanyu Hazarika, Roger Zimmermann, and Soujanya Poria. Misa: Modality-invariant and-specific representations for multimodal sentiment analysis. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1122–1131, 2020.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [23] Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, 2018.
- [24] Jian Huang, Jianhua Tao, Bin Liu, Zheng Lian, and Mingyue Niu. Multimodal transformer fusion for continuous emotion recognition. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3507–3511. IEEE, 2020.

- [25] Yash Kant, Dhruv Batra, Peter Anderson, Alexander Schwing, Devi Parikh, Jiasen Lu, and Harsh Agrawal. Spatially aware multimodal transformers for textvqa. In *16th European Conference on Computer Vision, ECCV 2020*, pages 715–732. Springer Science and Business Media Deutschland GmbH, 2020.
- [26] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *CoRR*, 2017.
- [27] Douwe Kiela, Suvrat Bhooshan, Hamed Firooz, Ethan Perez, and Davide Testuggine. Supervised multimodal bitransformers for classifying images and text. *arxiv:1909.02950*, 2020.
- [28] Y Alex Kolchinski and Christopher Potts. Representing social media users for sarcasm detection. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1115–1121, 2018.
- [29] Herbert M Lefcourt and Rod A Martin. *Humor and life stress: Antidote to adversity*. Springer Science & Business Media, 2012.
- [30] Lizhen Liu, Donghai Zhang, and Wei Song. Modeling sentiment association in discourse for humor recognition. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 586–591, 2018.
- [31] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [32] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2018.
- [33] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Viltbert: pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 13–23, 2019.
- [34] Grégoire Mialon, Dexiong Chen, Alexandre d’Aspremont, and Julien Mairal. A trainable optimal transport embedding for feature aggregation and its relationship to attention. In *ICLR 2021-The Ninth International Conference on Learning Representations*, 2021.
- [35] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [36] David Olson. From utterance to text: The bias of language in speech and writing. *Harvard educational review*, 47(3):257–281, 1977.
- [37] Hongliang Pan, Zheng Lin, Peng Fu, Yatao Qi, and Weiping Wang. Modeling intra and inter-modality incongruity for multi-modal sarcasm detection. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 1383–1392, 2020.
- [38] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019.
- [39] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- [40] Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, 2018.
- [41] Hai Pham, Paul Pu Liang, Thomas Manzini, Louis-Philippe Morency, and Barnabás Póczos. Found in translation: Learning robust joint representations by cyclic translations between modalities. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6892–6899, 2019.
- [42] Soujanya Poria, Erik Cambria, Rajiv Bajpai, and Amir Husain. A review of affective computing: From unimodal analysis to multimodal fusion. *Information Fusion*, 37:98–125, 2017.
- [43] Soujanya Poria, Iti Chaturvedi, Erik Cambria, and Amir Husain. Convolutional mkl based multimodal emotion recognition and sentiment analysis. In *2016 IEEE 16th international conference on data mining (ICDM)*, pages 439–448. IEEE, 2016.
- [44] Shraman Pramanick, Md Shad Akhtar, and Tanmoy Chakraborty. Exercise? I thought you said ‘extra fries’: Leveraging sentence demarcations and multi-hop attention for meme affect analysis. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, pages 513–524, 2021.
- [45] Shraman Pramanick, Dimitar Dimitrov, Rituparna Mukherjee, Shivam Sharma, Md Shad Akhtar, Preslav Nakov, and Tanmoy Chakraborty. Detecting harmful memes and their targets. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2783–2796, 2021.
- [46] Shraman Pramanick, Shivam Sharma, Dimitar Dimitrov, Md Shad Akhtar, Preslav Nakov, and Tanmoy Chakraborty. Momenta: A multimodal framework for detecting harmful memes and their targets. *arXiv preprint arXiv:2109.05184*, 2021.
- [47] Wasifur Rahman, Md Kamrul Hasan, Sangwu Lee, Amir Zadeh, Chengfeng Mao, Louis-Philippe Morency, and Ehsan Hoque. Integrating multimodal information in large pre-trained transformers. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2020, page 2359. NIH Public Access, 2020.
- [48] Ellen Riloff, Ashequl Qadir, Prafulla Surve, Lalindra De Silva, Nathan Gilbert, and Ruihong Huang. Sarcasm as contrast between a positive sentiment and negative situation. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 704–714, 2013.
- [49] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-

- cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [50] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypertexted, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018.
- [51] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7464–7473, 2019.
- [52] Zhongkai Sun, Prathusha Sarma, William Sethares, and Yingyu Liang. Learning relationships between text, audio, and video via deep canonical correlation for multimodal language analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8992–8999, 2020.
- [53] Joseph Tepperman, David Traum, and Shrikanth Narayanan. "yeah right": Sarcasm recognition for spoken dialogue systems. In *Ninth international conference on spoken language processing*, 2006.
- [54] Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J. Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6558–6569, Florence, Italy, July 2019. Association for Computational Linguistics.
- [55] Robert A Vartabedian and Laurel Klinger Vartabedian. *Humor in the workplace: A communication challenge*. 1993.
- [56] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [57] C Villani. Optimal transport, old and new. notes for the 2005 saint-flour summer school. *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer, 2008.
- [58] Yansen Wang, Ying Shen, Zhun Liu, Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. Words can shift: Dynamically adjusting word representations using nonverbal behaviors. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7216–7223, 2019.
- [59] Xi Wei, Tianzhu Zhang, Yan Li, Yongdong Zhang, and Feng Wu. Multi-modality cross attention network for image and sentence matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10941–10950, 2020.
- [60] Jennifer Woodland and Daniel Voyer. Context and intonation in the perception of sarcasm. *Metaphor and Symbol*, 26(3):227–239, 2011.
- [61] Nan Xu, Zhixiong Zeng, and Wenji Mao. Reasoning with multimodal sarcastic tweets via modeling cross-modality contrast and semantic association. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3777–3786, 2020.
- [62] Diyi Yang, Alon Lavie, Chris Dyer, and Eduard Hovy. Humor recognition and humor anchor extraction. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2367–2376, 2015.
- [63] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32, 2019.
- [64] Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Tensor fusion network for multimodal sentiment analysis. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1103–1114, 2017.
- [65] Amir Zadeh, Paul Pu Liang, Navonil Mazumder, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Memory fusion network for multi-view sequential learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [66] AmirAli Bagher Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2236–2246, 2018.