

Whose Emotion Matters? Speaker Detection without Prior Knowledge

Hugo Carneiro^a, Cornelius Weber^a, Stefan Wermter^a

^a*University of Hamburg, Department of Informatics, Vogt-Koelln-Str. 30, Hamburg, 22527, Germany*

Abstract

The task of emotion recognition in conversations (ERC) benefits from the availability of multiple modalities, as offered, for example, in the video-based MELD dataset. However, only a few research approaches use both acoustic and visual information from the MELD videos. There are two reasons for this: First, label-to-video alignments in MELD are noisy, making those videos an unreliable source of emotional speech data. Second, conversations can involve several people in the same scene, which requires the detection of the person speaking the utterance. In this paper we demonstrate that by using recent automatic speech recognition and active speaker detection models, we are able to realign the videos of MELD, and capture the facial expressions from uttering speakers in 96.92% of the utterances provided in MELD. Experiments with a self-supervised voice recognition model indicate that the realigned MELD videos more closely match the corresponding utterances offered in the dataset. Finally, we devise a model for emotion recognition in conversations trained on the face and audio information of the MELD realigned videos, which outperforms state-of-the-art models for ERC based on vision alone. This indicates that active speaker detection is indeed effective for extracting facial expressions from the uttering speakers, and that faces provide more informative visual cues than the visual features state-of-the-art models have been using so far.

Keywords:

multimodality, active speaker detection, emotion recognition, forced alignment

1. Introduction

Emotion recognition in conversations (ERC) is a task that involves recognising the emotion of interlocutors in a dialogue. Challenges of this task include the modelling of the conversational context and how the emotion of the interlocutors may change depending on that context, which is called emotion shift [31]. ERC can prove useful in real-world

Email addresses: hugo.carneiro@uni-hamburg.de (Hugo Carneiro), cornelius.weber@uni-hamburg.de (Cornelius Weber), stefan.wermter@uni-hamburg.de (Stefan Wermter)

scenarios in which people are talking with each other, for example, in human-robot interaction applications [17, 22, 39]. However, most ERC datasets are exclusively based on text transcriptions of conversations [18, 26, 42] or are restricted to dyadic interactions in very controlled environments [5, 29].

Recently, there has been an interest in large-scale multimodal ERC databases with several interlocutors, e.g., MELD (**M**ultimodal **E**motion**L**ines **D**ataset) [31], a dataset composed of videos extracted from the *Friends* TV series. Each video is cut to match a single utterance, and the videos are organised into dialogues and utterances, with each dialogue having one or more utterances. Together with the acoustic and visual information provided by the videos, the text transcription of every utterance and the speaker label are also provided.

Many approaches have been proposed to tackle the task of ERC in MELD. Even though MELD was created to be a multimodal dataset, most of the approaches rely exclusively on textual information [13, 24, 44, 25, 33, 36]. Using the visual modality is difficult due to frequent misalignments between video cuts and the expected corresponding utterances (see Figure 1 for an example). This is likely a consequence of an automatic generation of the video cuts with the *Gentle*¹ transcription alignment tool.

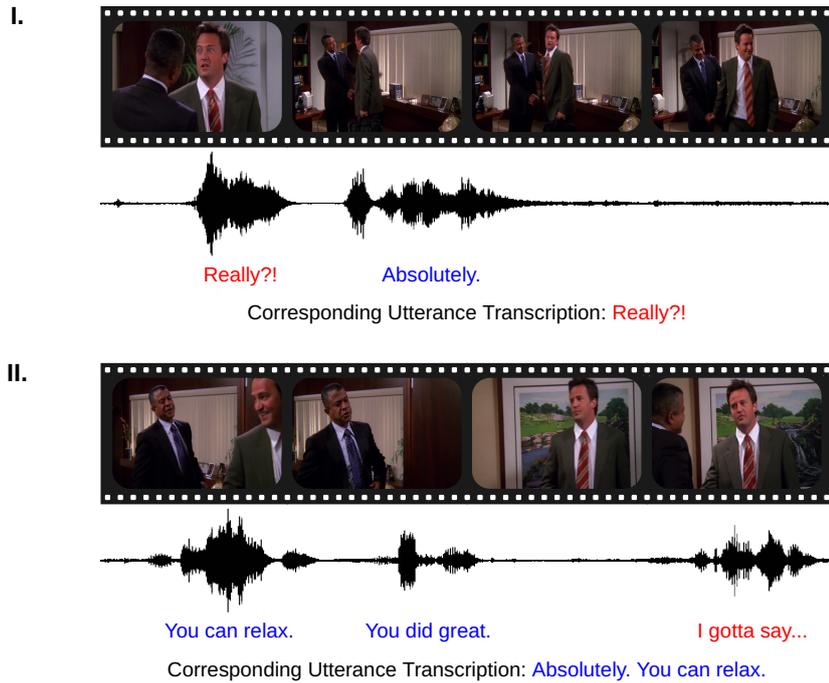
For some years there has been a demand for more reliable information from the visual modality given the frequent problems of video-text synchronisation². Video cuts and utterance transcriptions can be misaligned in a variety of ways. Figure 1a presents two cases of misalignment. In case I, the utterance appears within the first half of the video cut and another person’s utterance is falsely assigned to the same cut. In case II, the utterance starts being spoken in the video cut assigned to the preceding utterance, and continues through the first half of the video cut assigned to that target utterance. Figure 1b depicts the corrected alignment between the video cuts and their corresponding utterance transcriptions. This is a result of our dataset refinement procedure (cf. Section 3).

Facial expressions and speech signals provide relevant information regarding the emotion of a person. However, the noticeable number of mismatching cases between video cuts and the corresponding utterance transcriptions hindered the use of those modalities for some years, with information from the visual modality being disregarded even by the dataset creators, who stated that video-based speaker identification and localisation were still open problems [31]. Accordingly, in the dataset itself, no information on the location of the face of uttering speakers is offered.

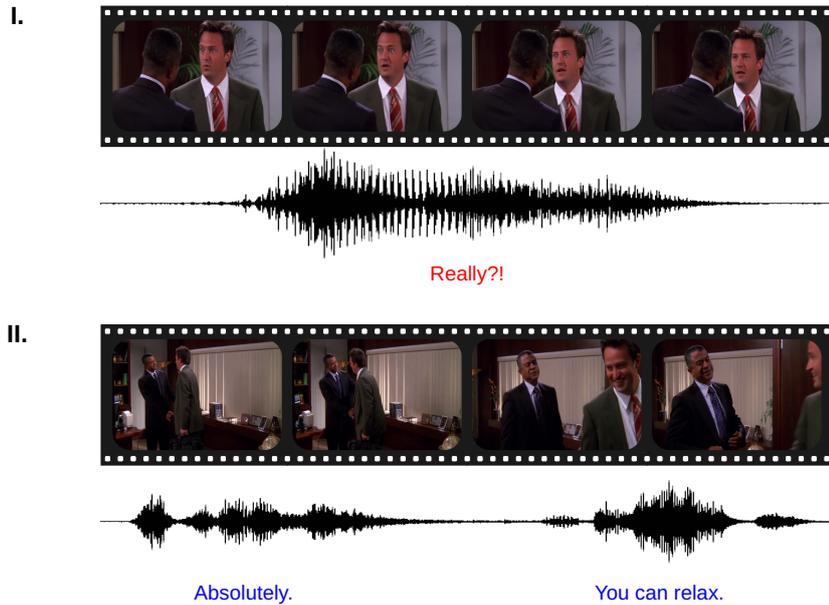
Even though quite rarely, speech data from the videos of MELD has been used for ERC since the work of Poria et al. [31]. While some preprocessing is done to the speech signals, no proper alignment correction has been made. Consequently, audio samples used for this task can include speech from other speakers with different emotions. The visual modality, in turn, was only used quite recently [41, 21, 19, 27, 9]. However, alongside the problems that arise with a lack of a proper realignment, the proposed solutions did not take into account the necessity to identify the utterance speaker in a particular scene or frame. The added information from acoustic and visual modalities improved ERC compared to models that

¹<https://lowerquality.com/gentle/>

²<https://github.com/declare-lab/MELD/issues/9>



(a) Original video cuts falsely corresponding to two consecutive utterances. The expected corresponding utterances are given below each video.



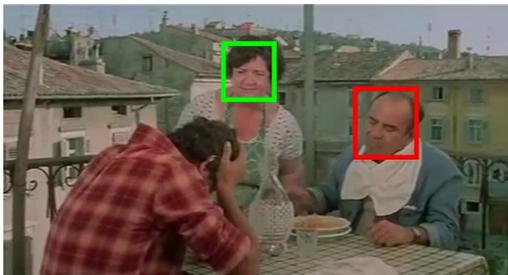
(b) Realigned video cuts, correctly matching the corresponding utterances.

Figure 1: Example of misaligned video cuts provided in the MELD dataset, and their corresponding correction. The different colours in the utterances represent the different speakers in the video cuts.

use information exclusively from the utterance transcriptions. However, those improvements are limited because of the unreliability of those modalities.

Recent advances in active speaker detection (ASD) in the wild [2, 3, 6, 30, 37, 43] indicate

the capability of audio-visual neural models to identify speaking activity in videos given the faces of the people as well as the audio of a scene. Identifying the active speaker’s face can enable a more reliable emotion recognition from video in MELD. State-of-the-art ASD models can be very precise in determining who among multiple people is speaking, especially if there are at least a few seconds of continuous speaking activity. Most in-the-wild ASD models were trained on AVA-ActiveSpeaker, a dataset containing videos in a large variety of resolutions [32]. The videos of AVA-ActiveSpeaker contain scenes with multiple people speaking with each other, which is similar to the conversational scenes in the videos of MELD. Figure 2 displays examples of conversational scenes present in the videos of AVA-ActiveSpeaker dataset.



(a) Three people interacting, with two of them having their faces visible and only one speaking.



(b) Scene with several people, with two of them speaking simultaneously.

Figure 2: Examples of conversational scenes from AVA-ActiveSpeaker videos. Green boxes identify those who are speaking whereas silent people are marked with a red box.

The first contribution of this paper is to offer a new method to extract the position of the faces of active speakers for datasets, which can be useful for tasks in which the facial information may provide additional relevant information, but, for some reason, the face position is not given in the dataset. The procedure can be used in any dataset with humans speaking without annotation concerning the visual modality, e.g., the position of the speaker’s face. The second contribution of this paper is the evaluation of this procedure on the MELD dataset. Finally, to assess the applicability of the extraction of the faces of active speakers for the task of ERC, we propose an emotion recognition model whose outstanding performance on the visual data indicates that the faces extracted from the active speakers indeed provide an informative visual cue for the task of ERC.

The paper is structured as follows. Section 2 offers a brief overview and some specific details on the MELD dataset. Section 3 describes the procedure of dataset refinement, which consists of correcting the alignment between video cuts and the corresponding utterances, and determining the position of the face of the uttering speaker in each frame of the newly produced video cut. Section 4 provides quantitative analyses on the resulting dataset, comparing it with the characteristics of the original dataset that were provided in Section 2. In that same section, experiments are also provided, as a means to evaluate how well the resulting dataset applies to the task of emotion recognition. Section 5 discusses the results.

2. The MELD Dataset

MELD contains scenes from various episodes of the *Friends* TV series. Those scenes are denoted as dialogues, and each dialogue is organised as a sequence of utterances. For every utterance, there is a corresponding dataset entry containing the speaker identity, the emotion and the sentiment of the speaker. The annotated emotion can be either one of Ekman’s universal emotions (*joy, sadness, fear, anger, disgust, and surprise*), or *neutral* if no particular emotion was noticed by the dataset annotators.

MELD is split into three sets, denoted *train, dev, and test*. Each data record in those splits contains the following information: the utterance being said, its speaker, the emotion perceived in that utterance, the corresponding sentiment, a dialogue identifier, an utterance identifier, the season and episode of *Friends* in which that scene happened, a time stamp determining where that scene starts, as well as one determining where that scene ends. For every split, a dataset record can be uniquely identified by its dialogue identifier and its utterance identifier.

Table 1: Excerpt of a dyadic conversation from MELD *train* split with corresponding speaker, emotion, and sentiment information

Dia	Utt	Utterance	Speaker	Emotion	Sentiment
D0	U5	Now you’ll be heading a whole division, so you’ll have a lot of duties.	Interviewer	neutral	neutral
D0	U6	I see.	Chandler	neutral	neutral
D0	U7	But there’ll be perhaps 30 people under you so you can dump a certain amount on them.	Interviewer	neutral	neutral
D0	U8	Good to know.	Chandler	neutral	neutral
D0	U9	We can go into detail.	Interviewer	neutral	neutral
D0	U10	No, don’t. I beg of you!	Chandler	fear	negative

Table 2: Additional information in MELD about the utterances presented in Table 1. Overlaps between video cuts due to mistaken determination of start and end times of an utterance are marked in **bold**

Dia	Utt	Season	Episode	Start time	End time
D0	U5	S8	E21	0:16:41.126	0:16:44.337
D0	U6	S8	E21	0:16:48.800	0:16:51.886
D0	U7	S8	E21	0:16:48.800	0:16:54.514
D0	U8	S8	E21	0:16:59.477	0:17:00.478
D0	U9	S8	E21	0:17:00.478	0:17:02.719
D0	U10	S8	E21	0:17:02.856	0:17:04.858

Table 1 presents an excerpt of a conversation, containing a sequence of contiguous data records and corresponding labels for the uttering speaker, his or her emotion, and the cor-

responding sentiment. The misalignment of the video cuts can provide overlaps, which are indicated by the start and end time stamps. Table 2 indicates that two videos of consecutive utterances present an overlap due to a wrongly executed alignment process.

3. Dataset Refinement Procedure

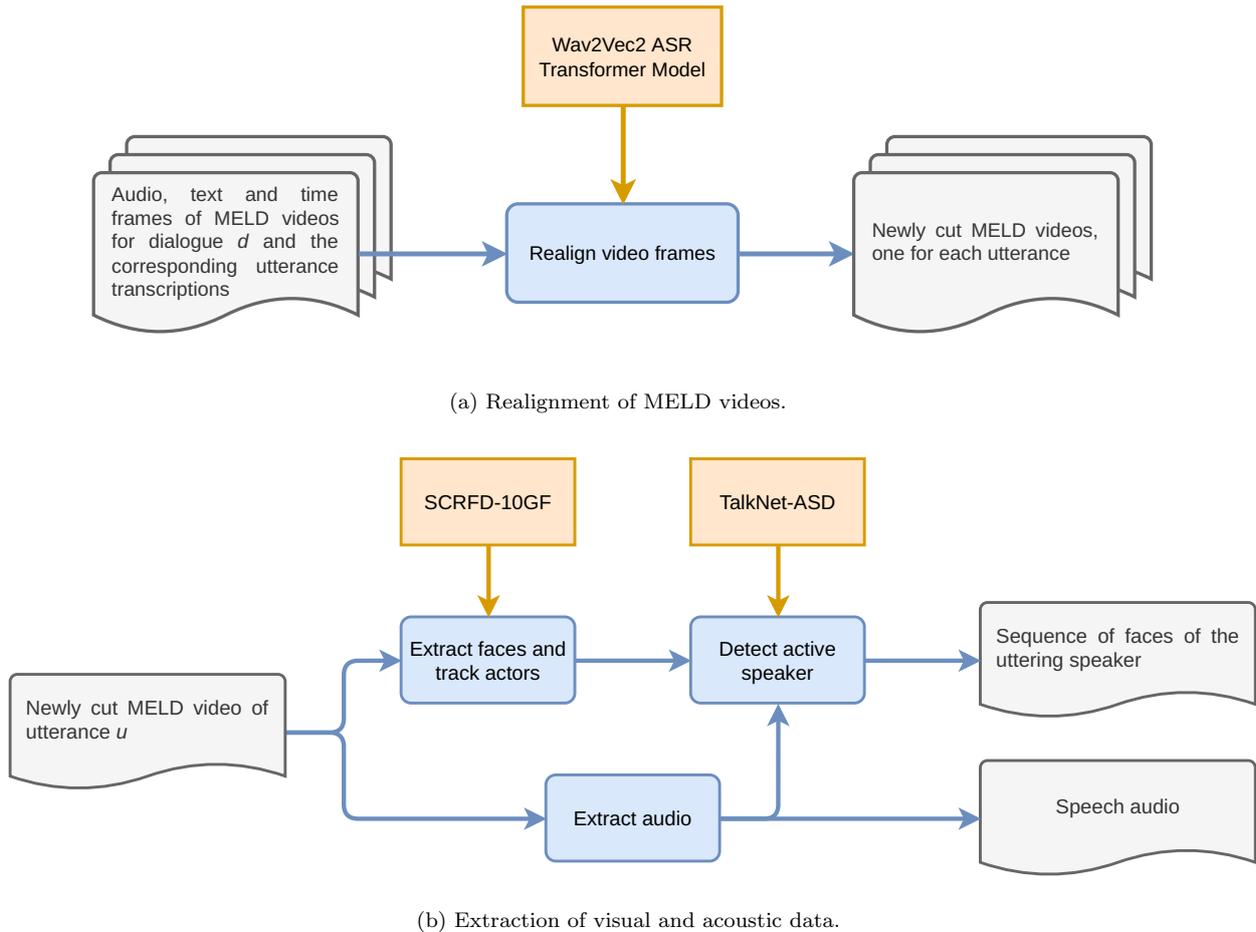


Figure 3: Major steps of the dataset refinement procedure. Orange arrows indicate application of a model, blue arrows represent information flow.

The extraction of emotional speech and emotional facial expressions depends on having audio samples that match closely enough the utterance being said and on being capable of determining the position of the uttering speaker in a scene, particularly that person’s face. For that, the dataset refinement procedure is divided into two parts. First, the videos of MELD are realigned, such that their audios match closely enough the target utterance, as indicated in the flowchart in Figure 3a. Next, with the videos properly realigned, the faces of the people in the scene are extracted and organised into sequences. Then, given the extracted sequences of faces and the scene audio, an ASD model determines which of those sequences corresponds to the uttering speaker (see the flowchart in Figure 3b).

3.1. Video Realignment

Each video in the MELD dataset corresponds to a particular utterance, which, in turn, belongs to a sequence of utterances, denoted a dialogue. Videos that are misaligned to their corresponding utterances are a consequence of the mistaken determination of where the boundaries of those particular utterances lie within their respective dialogues. A considerable number of misaligned videos prevents the proper identification of the uttering speakers, especially because sometimes the speaking activity might happen partially in the target video and partially in the one that precedes or in the one that follows it in the dialogue (see the example depicted in Figure 1).

The realignment of the videos takes into account that videos that belong to the same dialogue are organised in a sequential way. First, the audio signal of every dataset video is extracted. Next, for every split $\sigma \in \{train, dev, test\}$ and dialogue d , the audio signals $a_{\sigma,d,u}$ corresponding to each utterance u belonging to dialogue d are concatenated in order. Existing overlaps, as the one indicated in Table 2, are removed by truncating the audio signals that lead to those overlaps. Silence blocks are added between consecutive video cuts if there is a time difference between the end time stamp of a video and the start time stamp of the following. The length of a silence block is equal to the corresponding time difference, but long silence blocks are capped at 250 ms. Due to a few videos whose length is much longer than their corresponding utterances, video lengths are also capped to at most 45 seconds. This affects two of altogether 13,708 videos (check Appendix A for an indication of the videos affected by the 45-second capping).

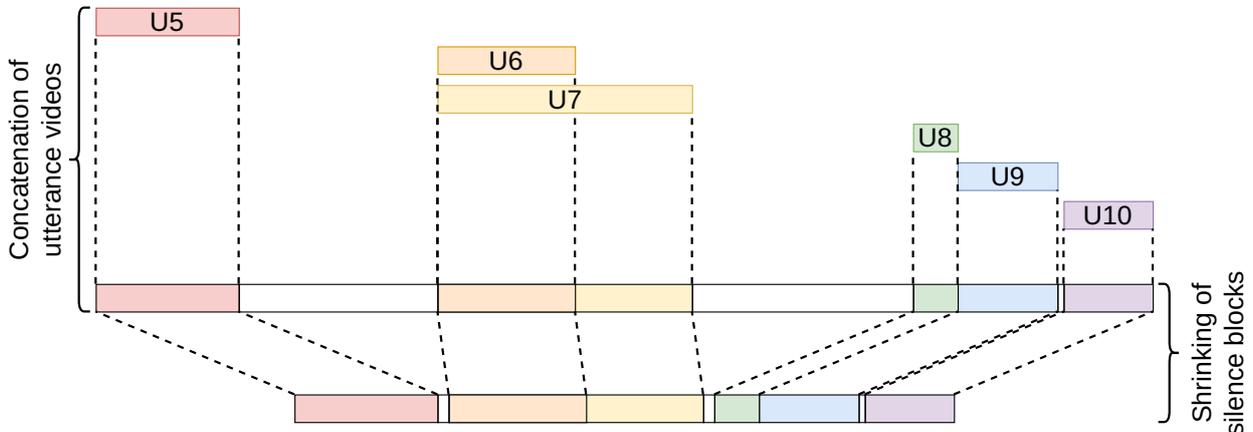


Figure 4: Schematic representation of the concatenation of the audio signals of a dialogue. First the audios of all utterances of a dialogue are concatenated, with silence blocks inserted wherever it is adequate. Next, the lengths of the silence blocks are reduced to a minimum length that still allows for the identification of individual blocks of consecutive utterances (e.g., utterances U6 and U7, and U8 and U9).

Figure 4 presents a graphical representation of the concatenation of audio signals. Each box labelled U5 to U10 represents the audio signal of an utterance. The lengths of those boxes are proportional to the difference between the end and start time stamps of the utterances presented in Tables 1 and 2. The label used in each box corresponds to the utterance identifier given in Table 2. The gaps between the boxes is also proportional to the

difference between the start time stamp of an utterance and the end time stamp of the ones that precedes it. The figure presents an example of overlap being removed by altering the start time of utterance U7. It also shows the insertion of silence blocks where lie the gaps between utterances, and the subsequent capping of silence block lengths to 250 ms.

The utterance transcriptions are concatenated as well. Prior to their concatenation, all punctuation marks in a transcription are removed, and both a start-of-sequence and an end-of-sequence token are appended to each end of every utterance transcription within a dialogue. With the audio signals and the transcriptions properly concatenated, the text of the concatenated transcription is aligned to the concatenated audio through a process of forced alignment using CTC (connectionist temporal classification) segmentation [23]. Given a speech audio signal, CTC segmentation uses frame-based character posterior probabilities generated by a CTC-based end-to-end network. From these character-level probabilities, maximum joint probabilities are computed via dynamic programming. These maximum joint probabilities indicate how likely a given excerpt from the dialogue transcription is aligned to a particular slice of the speech audio signal. After the maximum joint probability for the alignment of the complete dialogue transcription to the whole speech audio signal is computed, the character-wise alignment is obtained by backtracking from the most probable temporal position of the last character in the transcription. The CTC-based end-to-end network used to generate the character-level probabilities must have been pretrained on already aligned data beforehand, for which the Wav2Vec2 [4] automatic speech recognition transformer model³ [40] was used.

3.2. Uttering Speaker Determination

With videos that very likely contain the part of a scene in which a given utterance was said, it is possible to determine the source of the speaking activity, i.e., the person who spoke the utterance. Figure 3b schematically represents the process of extracting the speech audio as well as face images of the uttering speaker from a video. As a first step, an efficient face detection model with sample and computation redistribution (SCRFD-10GF) [15] was used to detect all faces in every frame of those videos. Faces detected this way were then subsequently extracted and organised into ordered groups, creating several sequences of faces. Each face is identified by the video frame from which it was extracted, and by an identifier of the sequence it belongs to.

Each face sequence and the corresponding slice of the speech signal was then sent to TalkNet-ASD [37], an audio-visual ASD model, to determine whether that face sequence presented some indication of a speaking activity that resembles that slice of the speech signal. Figure 5 shows a sketch of TalkNet-ASD’s architecture. TalkNet-ASD uses a visual temporal encoder (VTE) to learn long-term representations of facial expression dynamics, and an audio temporal encoder (ATE) to learn audio content representations from the temporal dynamics [37].

³More specifically, the Wav2Vec2 Large (LV-60) model pretrained and fine-tuned on 960 hours of speech audio from Libri-Light and Librispeech (see list of pretrained models at <https://github.com/facebookresearch/fairseq/tree/main/examples/wav2vec>).

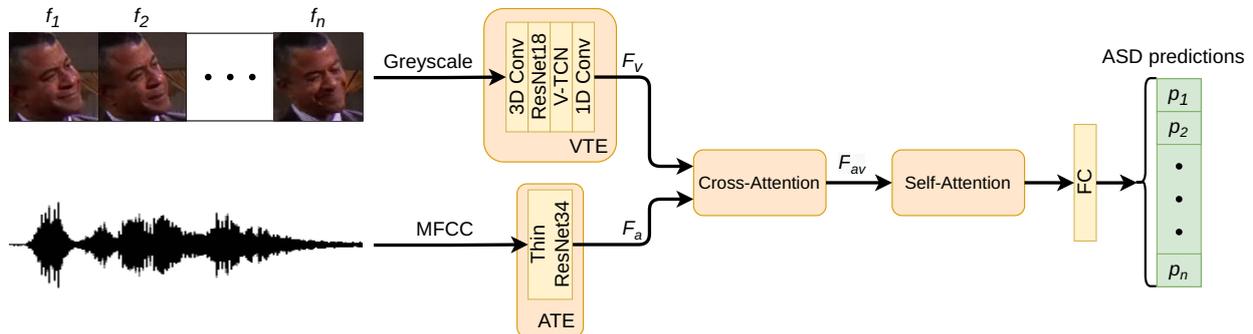


Figure 5: TalkNet-ASD architecture

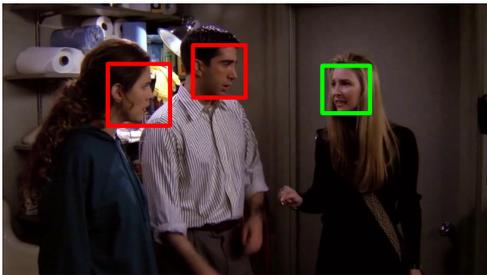
VTE consists of a front end, where video frame streams are encoded into sequences of frame-based embeddings, and a visual temporal network, whose aim is to represent the temporal content in a long-term spatio-temporal structure [37]. Its front end is based on the vision module introduced in [1], consisting on a 3D convolution layer with a filter width of 5 frames followed by a 2D 18-layer residual network. Given an input with dimensions $T_v \times C \times W \times H$, where T_v is the number of frames, and C , W and H are the number of channels, width and height of each frame, the front end yields a tensor with dimensions $T_v \times \frac{W}{32} \times \frac{H}{32} \times 512$, which is subsequently average-pooled in both its spatial dimensions, thus producing a feature vector with 512 dimensions for each input frame. Similarly to the visual model of Afouras et al. [1], TalkNet-ASD receives a sequence of greyscale images, which means that the number of channels C in each frame is 1. TalkNet-ASD’s visual temporal network (V-TCN) consists of a 5-block residual network followed by a sequence of two 1D convolution layers. The residual blocks consist of a 1D depth-separable convolution layer followed by rectifier linear units and batch normalisation layers. The residual network is responsible for obtaining a representation of the temporal content. The representation consists of a tensor with dimensions $T_v \times 512$. The sequence of 1D convolution layers, finally, reduces the dimensionality of this tensor, yielding a visual embedding F_v of dimensions $T_v \times 128$, i.e., 128 dimensions for every input frame.

The speech signal is first encoded as a sequence of overlapping audio frames, each one characterised by a 13-dimensional vector of Mel-frequency cepstral coefficients (MFCCs) based on an window size of 25 ms and a window step of 10 ms. This means that given a sequence of T_a audio frames, ATE receives as input a tensor with dimensions $1 \times 13 \times T_a$. ATE consists of a 2D 34-layer residual network with squeeze-and-excitation (SE) modules [20]. The number of channels in each block of the ResNet34 network is also reduced to only one quarter to those in each block of the original ResNet with 34 layers, similarly to the Thin ResNet34 introduced by Chung et al. [10]. The output of the audio encoder is an audio embedding F_a of dimensions $\frac{T_a}{4} \times 128$. The dimensions of F_a and F_v , the embeddings output by both encoders, match when the number of audio frames is equal to four times the number of video frames (or face crops). The matching in their dimensions is a necessary feature for the subsequent attention mechanism. A direct implication of having the number of audio frames being four times the number of video frames is that each video frame corresponds to

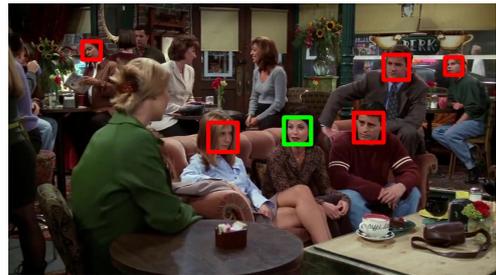
roughly 40 milliseconds of the video (or 25 fps), since the length of the window step between consecutive overlapping audio frames is 10 milliseconds.

With the motivation of audio-visual synchronisation working as an informative cue for speaking activities, TalkNet-ASD contains a cross-attention subnetwork that receives F_a and F_v as inputs and outputs an audio attention feature $F_{a \rightarrow v}$ and a video attention feature $F_{v \rightarrow a}$. $F_{a \rightarrow v}$ is obtained through the application of F_v as target sequence to generate the query Q_v in the attention layer and F_a as source sequence to generate key K_a and value V_a . $F_{v \rightarrow a}$ is obtained through an analogous process. Next, $F_{a \rightarrow v}$ and $F_{v \rightarrow a}$ are concatenated into a single audio-visual attention feature vector F_{av} , which is sent to a self-attention subnetwork, whose aim is to model audio-visual utterance-level information, and this way distinguish between speaking and non-speaking frames. Both cross-attention and self-attention subnetworks contain one transformer layer with eight attention heads [40].

Tao et al. [37] offer a practical implementation of TalkNet-ASD⁴, which we applied to the facial expression and emotional speech data extracted from the realigned MELD videos. In that implementation, each of the face tracks of a given person and the corresponding audio frame sequence are split into blocks and sent to TalkNet-ASD to determine in which frames that given person was actively speaking. Each of those blocks corresponds to a video sequence of up to ϕ video frames. Several values for ϕ were used in the implementation, namely 25, 50, 75, 100, 125, and 150, as a means to guarantee a more reliable result. A given value of ϕ implies that ϕ face images and 4ϕ audio frames in each block are used as input to the TalkNet-ASD model. TalkNet-ASD yields ϕ scores $s_{i,j,\phi}$ per block, indicating whether a given person π_j is detected as actively speaking in frame f_i in that block composed of ϕ video frames. After getting all scores for every frame, with all different possible values of ϕ , a resulting score $s_{i,j}$ is obtained by averaging the scores $s_{i,j,\phi}$. A score $s_{i,j} > 0$ indicates that person π_j is predicted as actively speaking in frame f_i . Figure 6 provides two examples of the application of TalkNet-ASD to the videos of MELD. In both examples, the uttering speakers are marked with green boxes around their faces. Figure 6 provides two examples of the application of TalkNet-ASD to the videos of MELD. In both examples, the uttering speakers are marked with green boxes around their faces.



(a) Conversation between three people in a low-lighting environment.



(b) Conversational scene in a crowded environment.

Figure 6: Examples of conversational scenes from MELD videos. Green boxes identify those who are speaking whereas silent people are marked with a red box.

⁴<https://github.com/TaoRuijie/TalkNet-ASD/blob/main/demoTalkNet.py>

4. Assessment of the Refined MELD Dataset

To assess the applicability of the refined dataset in ERC, it is important to determine whether the distribution of its data after the dataset refinement procedure is kept similar to that of the original dataset. Two criteria can be used to evaluate whether the data distribution was kept similar to its original distribution. Specific steps of the dataset refinement depend on the target uttering speaker, thus it is desirable that the proportion of utterances in the refined dataset assigned to a given speaker remains close to its original proportion. Similarly, the proportion of utterances assigned to a given emotion should also be kept close to its original proportion as a means to keep the task challenge compatible to its original difficulty. Moreover, because MELD was built for emotion recognition in conversational contexts, it is worthwhile to determine the portion of dialogues in which the data of at least one utterance was removed during the dataset refinement process.

If most of the original utterances are kept in the refined dataset, and its data distribution is nearly unaltered, then one should still analyse whether the video realignment produces refined speech signals that actually correspond to the speakers provided by the dataset. If the analysis indicates that the acoustic data is reliable and therefore useful for an application in ERC, i.e., many of the original utterances were kept, and the speech signals correspond well to the expected speakers, then a test case should be performed by comparing the performance of an emotion recognition model trained on the refined version of MELD with existing approaches to ERC in the original version of MELD that use visual and/or acoustic modalities. A high performance in this test case would also indicate the reliability of the process of determining the uttering speaker.

4.1. Properties of the Refined MELD Dataset

The process of dataset refinement consists of two steps, video realignment and uttering speaker detection. These refining steps may eventually lead to some utterances of the original dataset not having corresponding audio-visual data in the resulting refined dataset. This may happen due to two main reasons. First, the video realignment step may produce an empty video for a given utterance in case the CTC segmentation algorithm determines that in the most likely alignment, u_i is aligned to a very small slice of the dialogue audio. Second, even when new video cuts are produced in the video realignment step, no uttering speaker may be detected in the scene. Tables 3 and 4 present the number of dataset records for which there is corresponding audio-visual data in the refined dataset alongside the number of dataset records in its original version. Table 3 presents the dataset record distribution according to the annotated emotion and dataset split, and Table 4 presents the dataset record distribution according to the utterance speaker and dataset split. Tables 3 and 4 show that the *dev* and *test* splits each lost approximately 2.5% of their records in the dataset refinement process. Regarding the *train* split, data loss due to the dataset refinement was also relatively small, with the audio-visual data of the resulting refined dataset corresponding to 96.7% of the utterances of the original dataset.

The data distribution was kept nearly unaltered. For instance, the largest data distribution difference occurred in the fraction of dataset records assigned to the neutral emotion in

Table 3: Distribution of emotion annotations in the refined dataset. Number of original dataset records for each emotion and split are given inside parentheses.

Emotion	<i>train</i>	<i>dev</i>	<i>test</i>	Total
neutral	4537 (4710)	461 (470)	1226 (1256)	6224 (6436)
joy	1683 (1743)	160 (163)	389 (402)	2232 (2308)
surprise	1158 (1205)	140 (150)	270 (281)	1568 (1636)
sadness	670 (683)	109 (111)	207 (208)	986 (1002)
fear	261 (268)	39 (40)	49 (50)	349 (358)
anger	1082 (1109)	150 (153)	339 (345)	1571 (1607)
disgust	267 (271)	22 (22)	67 (68)	356 (361)
	9658 (9989)	1081 (1109)	2547 (2610)	13286 (13708)

Table 4: Distribution of uttering speakers in the refined dataset. Number of original dataset records for each speaker and split are given inside parentheses.

Speaker	<i>train</i>	<i>dev</i>	<i>test</i>	Total
Rachel	1392 (1435)	158 (164)	350 (356)	1900 (1955)
Monica	1253 (1299)	130 (137)	338 (346)	1721 (1782)
Phoebe	1269 (1321)	183 (185)	277 (291)	1729 (1797)
Joey	1456 (1509)	146 (149)	399 (411)	2001 (2069)
Chandler	1243 (1283)	100 (101)	374 (379)	1717 (1763)
Ross	1410 (1459)	211 (217)	368 (373)	1989 (2049)
others	1635 (1683)	153 (156)	441 (454)	2229 (2293)
	9658 (9989)	1081 (1109)	2547 (2610)	13286 (13708)

the *train* split. The dataset refinement procedure could not retrieve corresponding audio-visual data to 173 from the original 4710 dataset records of the *train* split which were assigned to the neutral emotion, i.e., 1.73%. These dataset records, which correspond to one utterance each, are well dispersed throughout the whole dataset. As a consequence, the fraction of dialogues which lost at least one of its utterances in the dataset refinement procedure is relatively higher. 222 of the 1038 dialogues of the *train* split contain at least one utterance with no corresponding audio-visual data in the refined dataset, which represents 21.4% of the dialogues in that split. For the *dev* and *test* splits, this reduction was lower. 19 of the 114 dialogues of the *dev* split, i.e., 16.7%, have utterances with no corresponding audio-visual data in the refined dataset, and for the *test* split, 49 of its 280 dialogues, i.e., 17.5%.

4.2. Assessment of the Video Realignment

Due to the lack of an annotation of the correct start and end time stamps of each utterance, a self-supervised form of assessing the robustness of the video realignment procedure was devised. A video correctly realigned to its corresponding utterance would have most of its audio content comprised of a speech signal uttered by the speaker annotated in the corresponding dataset record. This allows training a speaker identification model with the

speech signals of the realigned videos of the *train* split, so that it generalises and correctly identifies the speakers from speech signals of the realigned videos of the remaining splits. However, the model would require a given speaker to appear in a reasonable number of MELD records in all dataset splits, but only six speakers appear consistently throughout all MELD splits. They are the six main characters: Rachel, Monica, Phoebe, Joey, Chandler, and Ross. The remaining speakers appear rarely, indicating that it is highly unlikely that the speaker identification model should be capable of generalising well from their speech.

4.2.1. Model

A speaker identification model is used to assess whether the speech audio in a given realigned video actually matches the speaker annotated in the corresponding MELD record. The speaker identification model is composed of an encoder part followed by a classifier part. Based on TalkNet-ASD’s ATE, a traditional ResNet34 is used as the encoder. This encoder produces an embedding F_a of dimensions $\frac{T_a}{4} \times 512$, where T_a is the number of audio frames corresponding to the speech signal. Then, via a temporal max pooling, a 512-dimensional feature vector is obtained for the whole speech signal. Finally, a fully connected layer outputs a prediction regarding the expected speaker of the speech signal from this feature vector.

4.2.2. Data Augmentation

Following the steps of Tao et al. [37], negative sampling is used to augment the available speech data. In negative sampling augmentation, data is augmented by combining it with some other data within the same batch, hereafter denoted noise data. The noise data should share the same label of the original data, i.e., it is expected that both the original and the noise speech signal have been uttered by the same speaker. Through randomly selecting noise data that has those characteristics, an interference is made by combining the original and the noise audio tracks, thus coming up with a mixture of both. By benefitting from the in-domain noise and the interference speakers from the training set itself, this approach presents three advantages in comparison to traditional augmentation through the addition of white noise: *i*) the interference data is not artificially generated; *ii*) there is no need for data outside the training set for the audio augmentation; and *iii*) by using audio samples from the same speaker, the interference provided in the data augmentation accentuates the characteristics of that speaker’s voice.

With a 50% chance, an audio sample is selected to be augmented this way, which means that within a batch roughly half of its samples are augmented. Audio samples selected this way are either circularly padded or trimmed to match the size of the original audio sample. A single batch typically has audio samples of very different sizes. In order to let all audio samples in the same batch have the same size, they were either circularly padded or trimmed so that every audio sample in the same batch have a length equal to the average of the lengths of the original audio samples. This way, the model would be trained with samples with a reasonable size, and at least half of the samples of a batch would consist of unpadded continuous audio samples.

4.2.3. Training Procedure

To train the speaker identification model, audio tracks were randomly sampled, such that there would be roughly the same number of audio samples for each class (the six main characters). Batches of size 64 were used in the training. Audio samples were augmented according to the aforementioned procedure. The model was trained by minimising a cross-entropy loss function using an ADAM optimiser with an initial learning rate of $1e-4$, whose value was decreased in half every ten epochs. The training procedure was kept running until there was a sequence of 30 epochs with no improvement in the weight F1-score of the *dev* split.

4.2.4. Results and Analysis

Figure 7 presents the confusion matrices obtained when evaluating the speaker identification model in MELD’s *test* split. A comparison is presented on how well the speaker identification can generalise what it learned from each character’s voice from the data of the original dataset (Figure 7a) and from that of its refined version (Figure 7b). The speaker identification model subjected to the refined dataset achieved a weighted F1-score of 78.32 in that dataset’s *test* split, whereas the speaker identification model subjected to the original dataset could achieved a weighted F1-score of 67.07 in the corresponding *test* split. The confusion matrices and the weighted F1-scores indicate that the video realignment leads to cuts that better match the expected speaker, which, in turn, indicates that it is highly likely that the audio content of those cuts closely match the corresponding utterances whose transcriptions are given in the dataset.

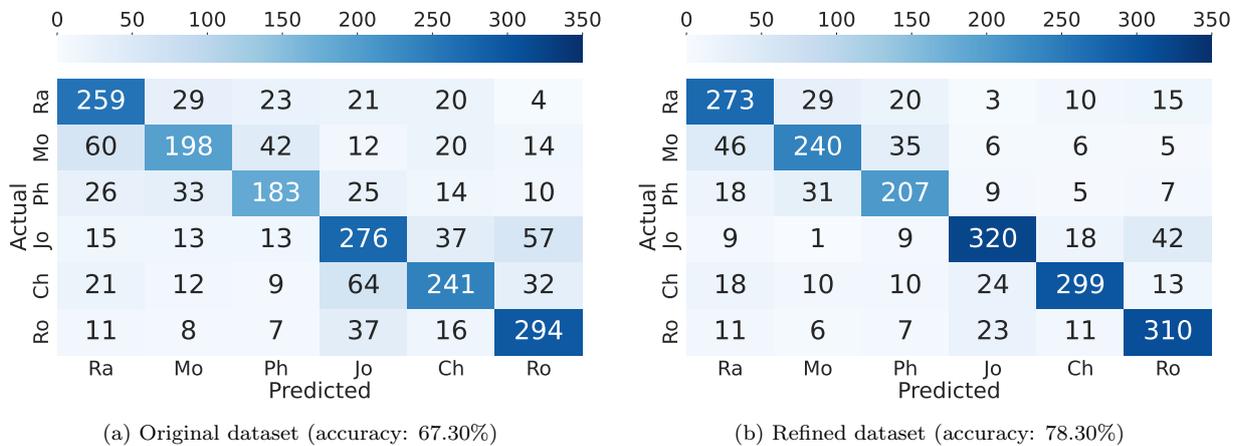


Figure 7: Confusion matrices of the speaker identification model in MELD *test* split

4.3. Application in ERC

4.3.1. Model

An emotion recognition (ER) model was devised to assess whether the refined version of MELD actually had visual and acoustic information from which emotional characteristics could be retrieved. Figure 8 presents the architecture of the ER model. For the encoding

of the visual and acoustic inputs, TalkNet-ASD’s VTE and ATE were used were modified for them to produce vector representations with 512 dimensions. For that, the sequence of 1D convolution layers of the VTE was removed, since their main application was to reduce the dimensionality of the feature vectors, and V-TCN already yielded vector representations with 512 dimensions. For TalkNet-ASD’s ATE to produce 512-dimensional feature vectors, its Thin ResNet34 backbone was changed for a traditional ResNet34. Also, we decided to keep the face crops with their original colour channels for the task of emotion recognition. This way, changes in skin colour due to some emotional reaction, e.g., blushing, could be considered by the ER model. The embeddings output by VTE and ATE are then max pooled in the temporal dimension into feature vectors F_v and F_a , with 512 dimensions each. These vectors are concatenated and subsequently sent to a self-attention layer. Finally, a fully connected layer yields a prediction for the emotion of the uttering speaker given the output of the self-attention layer.

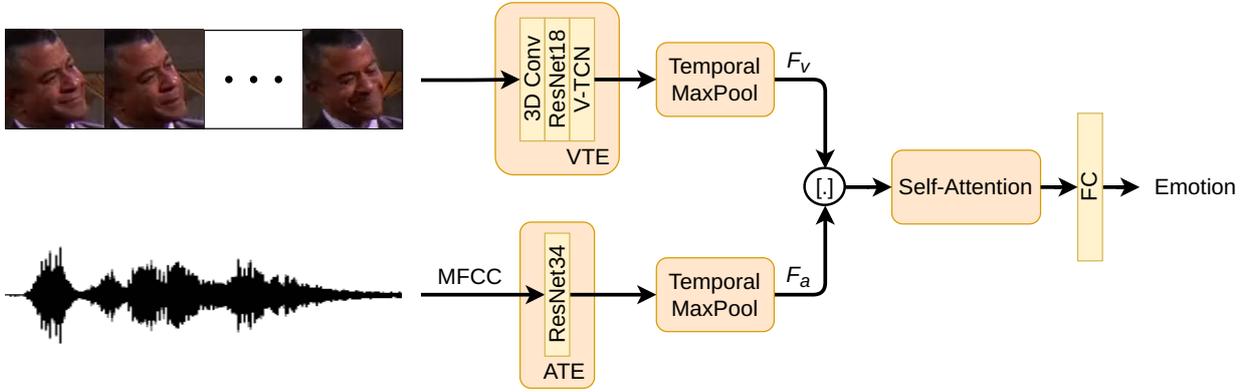


Figure 8: ERC model

4.3.2. Data Augmentation

Audio samples were augmented through the same data augmentation procedure described in Section 4.2. Face crops were augmented by performing one of the following operations: random horizontal flip, random crop of an area with at least 70% the dimension of the original face crop, or a random rotation up to 15 degrees clockwise or counterclockwise. Afterwards, the face crop was then again resized to 112×112 pixels. In order to keep consistency in the direction the head of speaker is looking to, the random characteristics of the data augmentation procedure were applied to the sequence of faces as a whole, and not to each face separately.

4.3.3. Training Procedure

Since the distribution of emotion labels is similar in every split of MELD, no weighted random sampling in the training of the ER model was performed. Instead, for every record in the *train* split representing a single utterance, a sequence of 15 consecutive face crops was selected as input for the video stream, and the complete utterance audio was provided as input for the audio stream. In case the sequence of faces corresponding to the uttering

speaker has less than 15 face crops, then the sequence is circularly padded. If the sequence of faces has more than 15 face crops, then a subsequence of 15 consecutive face crops is randomly selected. The model was trained by minimising a cross-entropy loss function. An ADAM optimiser was used, with an initial learning rate of 1e-4, whose value was decreased in half every ten epochs. Batches of size 64 were used in the training. The training procedure was kept running until there was a sequence of 30 epochs with no improvement in the weight F1-score of the *dev* split.

4.3.4. Comparison with the State-of-the-Art

To evaluate the benefits of our refinement procedure for the task of ERC with MELD, we compare the performance of our ER model to existing approaches that use information from the original MELD videos in ERC, and not only from the utterance transcriptions provided in the dataset.

DialogueRNN [28]⁵ is a baseline approach which models the context of a conversation by tracking the states of individual parties within that conversation. The model determines the emotion of a given utterance according to three aspects: its speaker, the context from preceding utterances, and the emotion thereof. DialogueRNN models these aspects by using three gated recurrent units (GRUs) [7], each responsible for a particular aspect.

CT+EmbraceNet [41] is a pioneering ERC model in using visual information from the MELD videos. Although DialogueRNN predates it, the former uses solely information from the acoustic and textual modalities. This approach uses crossmodal transformers (CTs) [38] to enrich the information from one modality by taking into account information from another modality, and this way learn existing correlated information across pairs of modalities. EmbraceNet [8] was used to carefully deal with the crossmodal information in the feature vectors produced by the crossmodal transformers, and to prevent performance degradation due to partial absence of data.

MMGCN [21] uses a multimodal graph where each node represents a given modality in some particular utterance. Nodes of this graph are connected if they share either the same modality or the same utterance. Each MMGCN node is initialised with a concatenation of a context-aware feature encoding of the corresponding modality and utterance and an embedding of the speaker of that particular utterance. MMGCN leverages the speaker embedding to inject speaker information into the graph construction. MMGCN encodes the multimodal contextual information through the use of a multilayered deep spectral-domain graph convolutional network.

MM-DFN [19], similarly to MMGCN, uses a multimodal graph with the same structure to characterise the relations between all modalities within a given uttering event, and of every utterance within a dialogue. MM-DFN introduces graph-based dynamic fusion modules, which are stacked in layers, to fuse multimodal context features dynamically and sequentially. These modules aggregate both inter- and intra-modality contextual information in a specific semantic space at each layer. It differs from MMGCN, which aggregates

⁵Although DialogueRNN was originally proposed in [28], its first application to ERC in the MELD dataset was in [31].

contextual information in a single semantic space, which leads to a gradual accumulation of redundant information. By modelling the contextual information in different semantic spaces, MM-DFN benefits from a reduction in the accumulation of redundant information, as well as from an enhancement in the complementarity between the modalities.

EmoCaps [27] uses transformer-based encoders to extract emotion feature vectors from the visual, acoustic and textual modalities. The authors also use BERT [11] to extract text feature vectors from every utterance. By concatenating an utterance feature vector with the corresponding emotion vectors of each modality, the authors create a vector representation for that utterance. Then, through the use of a Bi-LSTM [16, 14] and a classification subnetwork, EmoCaps predicts the emotion from every utterance in a dialogue.

M2FNet [9] is the current state-of-the-art model in ERC in MELD⁶. Its main characteristics are *i*) a visual feature extractor that provides a visual representation based on the faces of the people in a scene as well as on the scene as a whole; *ii*) the use of one stack of transformer encoders for each modality, as a means to learn inter-utterance context on a modality level; and *iii*) a multi-head attention fusion module to better incorporate those modalities, especially the visual and acoustic ones.

It is worth noticing that all multimodal approaches to ERC in MELD use context from the dialogue in some form. Since we are interested in extracting the most useful information from the visual and the acoustic modalities, we rely solely on the utterance level. This way, we can guarantee that the performance achieved is a direct consequence of the video realignment and the uttering speaker detection, and not from some other part of the dialogue.

Table 5: Weighted F1-scores for ERC in MELD *test* split using visual and acoustic data

Model	Vision	Audio	Audio + Vision
DialogueRNN	N/A	44.3	N/A
CT+EmbraceNet	31.4	32.1	N/A
MMGCN	33.27	42.63	N/A
MM-DFN	32.34	42.72	44.67
EmoCaps	31.26	31.26	N/A
M2FNet	32.44	39.63	35.74
Ours	35.58	40.54	39.81

Table 5 compares the performance of the ER model proposed here with those of ablated versions of all multimodal approaches to ERC in MELD. The values presented in the table were extracted from the literature. Some table cells appear empty because either one modality was not used (e.g., Poria et al. [31] do not use information from the visual modality in

⁶Although M2FNet’s performance values seem lower than those of other models in Table 5, this is due to most of the contribution in ERC coming from the textual modality, which was not included in Table 5. We decide not to include the performance of those models when the text modality is not ablated because the main objective of this paper is to present a way of extracting useful information from the visual and acoustic modalities, since those are quite unreliable in MELD. In contrast, the text transcriptions are very reliable and do not require an extensive refinement.

their implementation of DialogueRNN), or the authors did not consider the combination of vision and acoustic modalities in their ablation studies (as it was done in [21] and in [27]). Table 5 shows that our ER model achieves a higher weighted F1-score than state-of-the-art approaches when restricted to the vision modality. It is worth noticing that our ER model outperforms state-of-the-art approaches even though it does not use temporal visual context on a dialogue level. This indicates that the combination of video realignment and active speaker detection can indeed yield sequences of face expressions which, in turn, provide the ER model with more information on the uttering speaker’s emotion than the feature extraction procedures used in the other approaches.

The performance of our ER model when restricted to the acoustic modality is higher than M2FNet (current state-of-the-art approach for ERC in MELD) and EmoCaps. Its performance, however, is lower than those of DialogueRNN, MMGCN and MM-DFN. These models have in common the use of utterance-level feature vectors extracted from OpenSMILE [12, 35] as input for the audio stream. EmoCaps also uses these, however, its multimodal representation favours more the textual modality since it uses both the utterance feature vector yielded by BERT and an emotion feature vector for the textual modality in its multimodal utterance representation, whereas only a single emotion feature vector is used to represent each of the remaining modalities. Also, EmoCaps’s weighted F1-scores in both modalities correspond to that of a model that outputs *neutral* for every input. M2FNet, on the other hand, uses a novel feature extractor module based on the triplet loss [34] to fetch deep features from acoustic and visual contents.

5. Discussion and Conclusion

CTC segmentation and active speaker detection allowed us to refine MELD, making it possible to better align speech signals with their utterance transcriptions, as well as to obtain reliable face crops of the uttering speaker of nearly every scene. The comparison with state-of-the-art approaches also indicates that those face crops provide more precise information on the emotion of the uttering speaker than the most recent approaches. However, the performance obtained with emotion recognition using vision or speech alone does not take into account the conversational context.

The reliable extraction of the speakers’ face crops from well-realigned videos accounted for the high performance of the vision-only version of our ER model, whose weighted average F1-score is more than 2.3% higher than that of MMGCN. The relatively simple architecture of the ER model, as well as its restriction to working on an utterance-level, i.e., without contextual information from the whole dialogue, indicate that much of its high performance is due to the improvement in the information from MELD’s visual modality.

Furthermore, researchers on ERC in multi-party scenarios can benefit from the refined version of MELD delivered in this publication. More generally, with the recent advancements in deep learning, creating a dataset automatically becomes within sight. Automatic speech recognition allows automatic text transcription, while automatic lip reading, which requires active speaker detection, could verify its correctness, and vice versa. Meanwhile, the correctness of active speaker detection is verified by person identification from face and

from voice, which need to yield consistent identity. Person identification, in turn, requires training data, so we used the labels of MELD. If these do not exist, then some clustering approaches might be used instead, where face and voice would be required to map to the same cluster, without requiring any labels.

Acknowledgement

The authors acknowledge partial support from the German Research Foundation DFG under project CML (TRR 169).

Appendix A. Problematic cases of MELD

MELD presents a variety of problematic cases beyond the misalignment between the videos and the utterance transcriptions. These comprise multiple other problems which raised errors during the processing of data refinement. Table A.6 offers an extensive list of those cases, identified by the split, dialogue id and utterance id of each case.

Table A.6: List of existing problematic cases in MELD

Split	Dia ID	Utt ID	Problem
<i>train</i>	125	3	Corrupted video file
<i>dev</i>	110	7	Inexistent video file
<i>test</i>	38	4	Very long video (> 45 seconds), incompatible with its utterance transcription
<i>test</i>	220	0	
<i>train</i>	309	0	Utterance transcription contains not only the utterance but also a description within parentheses
<i>train</i>	404	15	
<i>train</i>	736	4	
<i>train</i>	832	3	
<i>train</i>	1018	2	
<i>dev</i>	108	0	
<i>test</i>	128	2	
<i>train</i>	65	3	Utterance transcription contains not only the utterance but also a description within brackets
<i>train</i>	761	1	
<i>test</i>	86	3	
<i>train</i>	739	14	No utterance. Just a description within parentheses
<i>train</i>	849	3	
<i>train</i>	111	N/A	Utterances not chronologically ordered
<i>train</i>	446	19	Should be the first utterance of dialogue 447

References

- [1] T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman. Deep audio-visual speech recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. doi: 10.1109/TPAMI.2018.2889052.

- [2] J. L. Alcázar, F. Caba, L. Mai, F. Perazzi, J.-Y. Lee, P. Arbelaez, and B. Ghanem. Active speakers in context. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 12465–12474, June 2020.
- [3] J. L. Alcázar, F. Caba, A. K. Thabet, and B. Ghanem. MAAS: Multi-modal assignment for active speaker detection. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pages 265–274, Oct. 2021.
- [4] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli. Wav2Vec 2.0: A framework for self-supervised learning of speech representations. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, Advances in Neural Information Processing Systems, volume 33, pages 12449–12460. Curran Associates, Inc., 2020.
- [5] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan. IEMOCAP: interactive emotional dyadic motion capture database. Language Resources and Evaluation, 42:335–359, 2008. doi: 10.1007/s10579-008-9076-6.
- [6] H. Carneiro, C. Weber, and S. Wermter. FaVoA: Face-voice association favours ambiguous speaker detection. In I. Farkaš, P. Masulli, S. Otte, and S. Wermter, editors, Artificial Neural Networks and Machine Learning – ICANN 2021, pages 439–450, Cham, 2021. Springer International Publishing.
- [7] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio. On the properties of neural machine translation: Encoder-decoder approaches. In Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation, pages 103–111, Doha, Qatar, Oct. 2014. Association for Computational Linguistics. doi: 10.3115/v1/W14-4012.
- [8] J.-H. Choi and J.-S. Lee. EmbraceNet: A robust deep learning architecture for multimodal classification. Information Fusion, 51:259–270, 2019. doi: 10.1016/j.inffus.2019.02.010.
- [9] V. Chudasama, P. Kar, A. Gudmalwar, N. Shah, P. Wasnik, and N. Onoe. M2FNet: Multi-modal fusion network for emotion recognition in conversation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, pages 4652–4661, June 2022.
- [10] J. S. Chung, J. Huh, and S. Mun. Delving into VoxCeleb: Environment invariant speaker recognition. In K. Lee, T. Koshinaka, and K. Shinoda, editors, Odyssey 2020: The Speaker and Language Recognition Workshop, 1-5 November 2020, Tokyo, Japan, pages 349–356. ISCA, 2020. doi: 10.21437/Odyssey.2020-49.
- [11] J. Devlin, M. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In J. Burstein, C. Doran, and T. Solorio, editors, Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers), pages 4171–4186. Association for Computational Linguistics, 2019. doi: 10.18653/v1/n19-1423.
- [12] F. Eyben, M. Wöllmer, and B. Schuller. OpenSMILE: The Munich versatile and fast open-source audio feature extractor. In Proceedings of the 18th ACM International Conference on Multimedia, MM ’10, page 1459–1462, New York, NY, USA, 2010. Association for Computing Machinery. doi: 10.1145/1873951.1874246.
- [13] D. Ghosal, N. Majumder, A. Gelbukh, R. Mihalcea, and S. Poria. COSMIC: COmmonSense knowledge for eMotion Identification in Conversations. In Findings of the Association for Computational Linguistics: EMNLP 2020, pages 2470–2481, Online, Nov. 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.224.
- [14] A. Graves and J. Schmidhuber. Framewise phoneme classification with bidirectional LSTM and other neural network architectures. Neural Networks, 18(5):602–610, 2005. doi: 10.1016/j.neunet.2005.06.042.
- [15] J. Guo, J. Deng, A. Lattas, and S. Zafeiriou. Sample and computation redistribution for efficient face detection. In International Conference on Learning Representations, pages 1–17, 2022. URL <https://openreview.net/forum?id=RhB1AdoFfGE>.
- [16] S. Hochreiter and J. Schmidhuber. Long short-term memory. Neural Computation, 9(8):1735–1780, Nov. 1997. doi: 10.1162/neco.1997.9.8.1735.
- [17] E. Hornecker, A. Krummheuer, A. Bischof, and M. Rehm. Beyond dyadic HRI: Building robots for

- society. *Interactions*, 29(3):48–53, May 2022. doi: 10.1145/3526119.
- [18] C.-C. Hsu, S.-Y. Chen, C.-C. Kuo, T.-H. Huang, and L.-W. Ku. EmotionLines: An emotion corpus of multi-party conversations. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, pages 1597–1601, Miyazaki, Japan, May 2018. European Language Resources Association (ELRA).
- [19] D. Hu, X. Hou, L. Wei, L. Jiang, and Y. Mo. MM-DFN: Multimodal dynamic fusion network for emotion recognition in conversations. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7037–7041, 2022. doi: 10.1109/ICASSP43922.2022.9747397.
- [20] J. Hu, L. Shen, and G. Sun. Squeeze-and-excitation networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7132–7141, 2018. doi: 10.1109/CVPR.2018.00745.
- [21] J. Hu, Y. Liu, J. Zhao, and Q. Jin. MMGCN: Multimodal fusion via deep graph convolution network for emotion recognition in conversation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5666–5675, Online, Aug. 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.440.
- [22] A. L. Krummheuer, M. Rehm, and K. Rodil. Triadic human-robot interaction. Distributed agency and memory in robot assisted interactions. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction, HRI '20*, page 317–319, New York, NY, USA, 2020. Association for Computing Machinery. doi: 10.1145/3371382.3378269.
- [23] L. Kürzinger, D. Winkelbauer, L. Li, T. Watzel, and G. Rigoll. Ctc-segmentation of large corpora for german end-to-end speech recognition. In A. Karpov and R. Potapova, editors, *Speech and Computer*, pages 267–278, Cham, 2020. Springer International Publishing. doi: 10.1007/978-3-030-60276-5_27.
- [24] B. Lee and Y. S. Choi. Graph based network with contextualized representations of turns in dialogue. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 443–455, Online and Punta Cana, Dominican Republic, Nov. 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.36.
- [25] J. Lee and W. Lee. CoMPM: Context modeling with speaker’s pre-trained memory tracking for emotion recognition in conversation. In M. Carpuat, M. de Marneffe, and I. V. M. Ruíz, editors, *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 5669–5679. Association for Computational Linguistics, 2022. doi: 10.18653/v1/2022.naacl-main.416.
- [26] Y. Li, H. Su, X. Shen, W. Li, Z. Cao, and S. Niu. DailyDialog: A manually labelled multi-turn dialogue dataset. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan, Nov. 2017. Asian Federation of Natural Language Processing.
- [27] Z. Li, F. Tang, M. Zhao, and Y. Zhu. EmoCaps: Emotion capsule based model for conversational emotion recognition. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1610–1618, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-acl.126.
- [28] N. Majumder, S. Poria, D. Hazarika, R. Mihalcea, A. Gelbukh, and E. Cambria. DialogueRNN: An attentive RNN for emotion detection in conversations. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):6818–6825, Jul 2019. doi: 10.1609/aaai.v33i01.33016818.
- [29] G. McKeown, M. Valstar, R. Cowie, M. Pantic, and M. Schroder. The SEMAINE database: Annotated multimodal records of emotionally colored conversations between a person and a limited agent. *IEEE Transactions on Affective Computing*, 3(1):5–17, 2012. doi: 10.1109/T-AFFC.2011.20.
- [30] K. Min, S. Roy, S. Tripathi, T. Guha, and S. Majumdar. Learning long-term spatial-temporal graphs for active speaker detection. In S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, and T. Hassner, editors, *Computer Vision – ECCV 2022*, pages 371–387, Cham, 2022. Springer Nature Switzerland. doi: 10.1007/978-3-031-19833-5_22.
- [31] S. Poria, D. Hazarika, N. Majumder, G. Naik, E. Cambria, and R. Mihalcea. MELD: A multimodal

- multi-party dataset for emotion recognition in conversations. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pages 527–536, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1050.
- [32] J. Roth, S. Chaudhuri, O. Klejch, R. Marvin, A. Gallagher, L. Kaver, S. Ramaswamy, A. Stopczynski, C. Schmid, Z. Xi, and C. Pantofaru. AVA active speaker: An audio-visual dataset for active speaker detection. In ICASSP 2020 – 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 4492–4496, 2020. doi: 10.1109/ICASSP40776.2020.9053900.
- [33] P. Saxena, Y. J. Huang, and S. Kurohashi. Static and dynamic speaker modeling based on graph neural network for emotion recognition in conversation. In Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Student Research Workshop, pages 247–253, Hybrid: Seattle, Washington + Online, July 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.naacl-srw.31.
- [34] F. Schroff, D. Kalenichenko, and J. Philbin. FaceNet: A unified embedding for face recognition and clustering. In 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 815–823, 2015. doi: 10.1109/CVPR.2015.7298682.
- [35] B. Schuller, A. Batliner, S. Steidl, and D. Seppi. Recognising realistic emotions and affect in speech: State of the art and lessons learnt from the first challenge. Speech Communication, 53(9):1062–1087, 2011. doi: 10.1016/j.specom.2011.01.011.
- [36] X. Song, L. Zang, R. Zhang, S. Hu, and L. Huang. Emotionflow: Capture the dialogue level emotion transitions. In ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 8542–8546, 2022. doi: 10.1109/ICASSP43922.2022.9746464.
- [37] R. Tao, Z. Pan, R. K. Das, X. Qian, M. Z. Shou, and H. Li. Is someone speaking? Exploring long-term temporal features for audio-visual active speaker detection. In Proceedings of the 29th ACM International Conference on Multimedia, MM ’21, page 3927–3935, New York, NY, USA, 2021. Association for Computing Machinery. doi: 10.1145/3474085.3475587.
- [38] Y.-H. H. Tsai, S. Bai, P. P. Liang, J. Z. Kolter, L.-P. Morency, and R. 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. doi: 10.18653/v1/P19-1656.
- [39] D. Utami and T. Bickmore. Collaborative user responses in multiparty interaction with a couples counselor robot. In 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI), pages 294–303, 2019. doi: 10.1109/HRI.2019.8673177.
- [40] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17, pages 6000–6010, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- [41] B. Xie, M. Sidulova, and C. H. Park. Robust multimodal emotion recognition from conversation with transformer-based crossmodality fusion. Sensors, 21(14), 2021. doi: 10.3390/s21144913.
- [42] S. Zahiri and J. D. Choi. Emotion detection on TV show transcripts with sequence-based convolutional neural networks. In Proceedings of the AAAI Workshop on Affective Content Analysis, AFFCON’18, pages 44–51, New Orleans, LA, 2018.
- [43] Y. Zhang, S. Liang, S. Yang, X. Liu, Z. Wu, S. Shan, and X. Chen. UniCon: Unified context network for robust active speaker detection. In Proceedings of the 29th ACM International Conference on Multimedia, MM ’21, pages 3964–3972, New York, NY, USA, 2021. Association for Computing Machinery. doi: 10.1145/3474085.3475275.
- [44] L. Zhu, G. Pergola, L. Gui, D. Zhou, and Y. He. Topic-driven and knowledge-aware transformer for dialogue emotion detection. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 1571–1582, Online, Aug. 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.125.