

# Continuous Emotion Recognition using Visual-audio-linguistic Information: A Technical Report for ABAW3

Su Zhang<sup>1</sup>, Ruyi An<sup>1</sup>, Yi Ding<sup>1</sup>, Cuntai Guan<sup>1,\*</sup>

<sup>1</sup>Nanyang Technological University

sorazcn@gmail.com, ran003@e.ntu.edu.sg, ding.yi@ntu.edu.sg, ctguan@ntu.edu.sg

## Abstract

*We propose a cross-modal co-attention model for continuous emotion recognition using visual-audio-linguistic information. The model consists of four blocks. The visual, audio, and linguistic blocks are used to learn the spatial-temporal features of the multi-modal input. A co-attention block is designed to fuse the learned features with the multi-head co-attention mechanism. The visual encoding from the visual block is concatenated with the attention feature to emphasize the visual information. To make full use of the data and alleviate over-fitting, cross-validation is carried out on the training and validation set. The concordance correlation coefficient (CCC) centering is used to merge the results from each fold. The achieved CCC on the test set is 0.520 for valence and 0.602 for arousal, which significantly outperforms the baseline method with the corresponding C-CC of 0.180 and 0.170 for valence and arousal, respectively. The code is available at <https://github.com/sucv/ABAW3>.*

## 1. Introduction

Emotion recognition is the process of identifying human emotion. It plays a crucial role in behavioral modeling, human-computer interaction, and affective computing. By using the dimensional model [1], any emotional state can be taken as a point located in a continuous space, with the valence and arousal being the axes. Continuous emotion recognition seeks to map the  $N$  sequential data points into  $M$  sequential emotional state points, where  $M$  usually equals  $N$ . This report details our methodology for the valence-arousal estimation challenge from the third affective behavior analysis in-the-wild (ABAW3) workshop [2–10].

Our work is an extension of the last year’s attempt [11] on ABAW2 [3]. The audio-visual-linguistic information is

exploited and combined by using a multi-head attentive feature fusion scheme. The network consists of three branches, fed by synchronous video frames, VGGish [12] features, and BERT features. The visual branch consists of a Resnet50 for spatial encoding and a temporal convolutional network (TCN) [13] for temporal encoding. The audio and linguistic branches each contain a TCN for temporal encoding. The three branches work in parallel and their outputs are sent to the attentive fusion block. The latter employs the co-attention mechanism of the three branches for feature fusion. To emphasize the dominance of the visual features, which we believe to have the strongest correlation with the label, the temporal feature outputted by the TCN of the visual branch is concatenated to the fused feature. A fully-connected layer is used for the regression. The 6-fold cross-validation is carried out to alleviate the over-fitting.

The remainder of the paper is arranged as follows. Section 2 details the model architecture including the visual, aural, linguistic, and attentive fusion blocks. Section 3 elaborates the implementation details including the data pre-processing, training settings, and post-processing. Section 4 provides the continuous emotion recognition results on the Aff-Wild2 database. Section 5 concludes the work.

## 2. Model Architecture

The model architecture is illustrated in Fig. 1. The visual block consists of a Resnet50 and a TCN. The resnet50 plays the role of backbone and is pre-trained on the MS-CELEB-1M dataset as a facial recognition task, it is then fine-tuned on the FER+ [14] dataset. The Resnet50 spatially encodes each frame from a video frame sequence. The sequential spatial encodings are then stacked and fed to a TCN for temporal encoding. The audio and linguistic blocks each consist of a TCN. The Vggish [12] and BERT [15] features are the inputs, respectively.

Given the three temporal encodings from the visual, audio, and linguistic blocks, the attentive fusion block first maps the feature vectors to query, key, and value vectors by the following procedure. For the  $i$ -th branch, its encoder

\*This work is partially supported by the RIE2020 AME Programmatic Fund, Singapore (No. A20G8b0102).

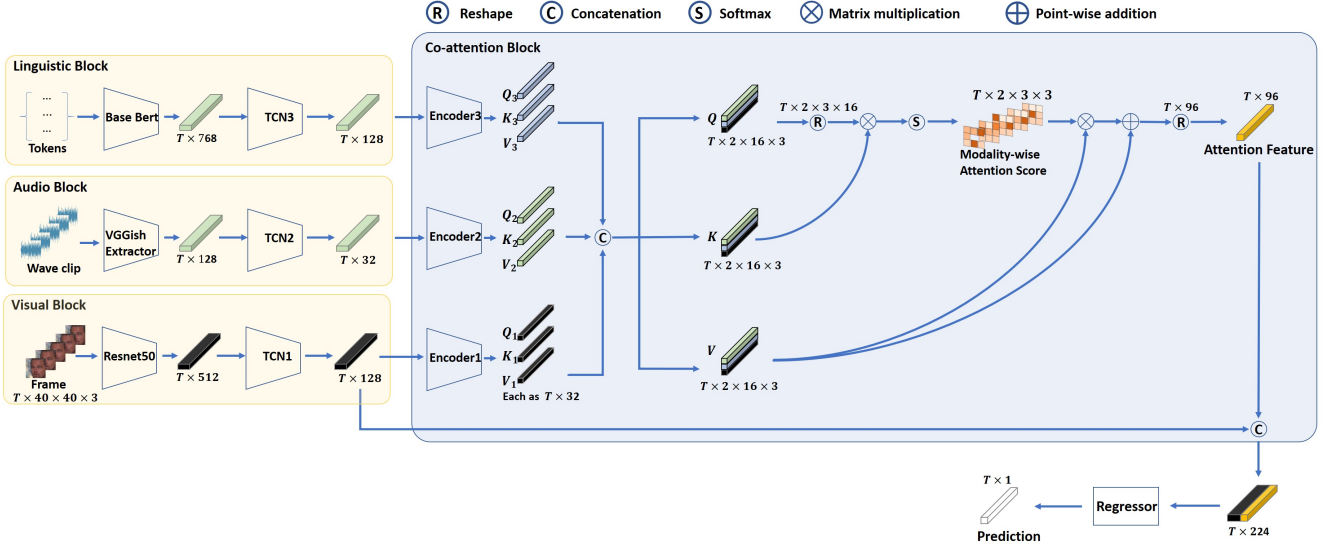


Figure 1. The architecture of our proposed model. The model consists of four components, i.e., the visual, audio, linguistic, and co-attention blocks. The visual block has a cascaded 2DCNN-TCN structure, and the audio and linguistic blocks each contain a TCN. The three branches yield three independent spatial-temporal feature vectors. They are then fed to the attentive fusion block. Three independent attention encoders are used. For the  $i$ -th branch, its encoder consists of three independent linear layers, they adjust the dimension of the feature vector producing a query  $\mathbf{Q}_i$ , a key  $\mathbf{K}_i$ , and a value  $\mathbf{V}_i$ . They are then regrouped and concatenated to form the cross-modal counterparts. For example, the cross-modal query  $\mathbf{Q} = [\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3]$ . An attention score is obtained by Eq. 1.

consists of three independent linear layers, they adjust the dimension of the feature vector producing a query  $\mathbf{Q}_i$ , a key  $\mathbf{K}_i$ , and a value  $\mathbf{V}_i$ . They are then regrouped and concatenated to form the cross-modal counterparts. For example, the cross-modal query  $\mathbf{Q} = [\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3]$ . After which, the attention feature is calculated as:

$$Attention(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = (\text{softmax}(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_K}}) + 1)\mathbf{V}, \quad (1)$$

where  $d_K = 32$  is the dimension of the key  $\mathbf{K}$ . After which, the attention feature is normalized and concatenated to the visual temporal encoding. Finally, a fully connected layer is employed to infer. Following the labeling scheme of the Aff-Wild2 database where the label frequency equals the video frame rate, each frame of the input video sequence is exactly corresponding to one label point.

### 3. Implementation Details

#### 3.1. Database

The ABAW3 competition uses the Aff-Wild2 database. The corpora of the valence-arousal estimation sub-challenge includes 564 trials. The database is split into the training, validation and test sets. The partitioning is done in a subject independent manner so that any subject's data are included in only one partition. The partitioning produces 341, 71, and 152 trials for the training, validation, and test

sets. Four experts annotate the videos using the method proposed in [16]. In addition to the annotations (for the training and validation sets only) and the raw videos, the bounding boxes and landmarks for each raw video are also available to the participants.

#### 3.2. Preprocessing

The visual preprocessing is carried out as follows. The cropped-aligned image data provided by the organizer are used. All the images are resized to  $48 \times 48 \times 3$ . Given a trial from the training or validation set, the length  $N$  is determined by the number of the rows of the annotation text file which does not include  $-5$ . For the test set, the length  $N$  is determined by the frame number of the raw video. A zero matrix  $\mathbf{B}$  of size  $N \times 48 \times 48 \times 3$  is initialized and then iterated over the rows. For the  $i$ -th row of  $\mathbf{B}$ , it is assigned as the  $i$ -th jpg image if it exists, otherwise doing nothing.

The audio preprocessing firstly converts all the videos to mono with a  $16K$  sampling rate in wav format. The VGGish features are then extracted using the pretrained Vggish model<sup>1</sup>. The only change is that we specified the hop length to be  $1/\text{framerate}$  of the raw video, in order to synchronize with other modalities and annotations.

The linguistic preprocessing is carried out as follows. The mono wav file obtained from the audio preprocessing is fed to a pretrained speech recognition model from the Vosk

<sup>1</sup><https://github.com/harritaylor/torchvggish>

toolkit<sup>2</sup>, from which the recognized words and the word-level timestamp are obtained. The recognized words are then fed to a pretrained punctuation and capitalization model from the Nvidia Nemo toolkit<sup>3</sup>. After which a pretrained BERT model from the Pytorch library is employed to extract the word-level linguistic features. The linguistic features are obtained by summing together the last four layers of the BERT model [17]. To synchronize, the word-level linguistic features are populated according to the timestamp of each word and each frame. Specifically, a word usually has a larger time span than that for a frame. Therefore, for one word, its feature is repetitively assigned to the time spans of all the frames within the time span.

For the valence-arousal labels, all the rows containing  $-5$  are excluded. To ensure that the features have the same length as the corresponding trial, the feature matrices are either repeatedly padded (using the last feature points) or trimmed (starting from the rear), depending on whether the feature length is shorter than the trial length or not, respectively.

### 3.3. Data Expansion

The AffWild2 database employed by ABAW3 contains 341 and 71 trials in the training and validation sets, respectively. To clarify, a label txt file and its corresponding data are taken as a trial. Note that some videos include two subjects, resulting in two separated cropped-aligned image folders and label txt files, with different suffixes. They are each taken as two trials. Note that there are 10 fewer trials for the training set compared to its counterpart from ABAW2.

To make full use of the available data and alleviate overfitting, 6-fold cross-validation is employed. By evenly splitting the training set into 5 folds, we have 6 folds in total with a roughly equal trial amount, i.e.,  $68 \times 4 + 69 + 71$  trials. Note that the 0-th fold is exactly the original data partitioning. And there is no subject overlap across different folds. The CCC-centering is employed to merge the inference result on the test set.

Moreover, during training and validation, the resampling window has a 33% overlap, resulting in 33% more data.

### 3.4. Training

The batch size is 2. For each batch, the resampling window length and hop length are 300 and 200, respectively. I.e., the dataloader loads consecutive 300 feature points to form a minibatch, with a stride of 200. For any trials having feature points smaller than the window length, zero padding is employed. For visual data, the random flip, random crop with a size of 40 are employed for training and only the

center crop is employed for validation. The data are then normalized so that  $mean = std = 0.5$ . For audio and linguistic data, they are normalized so that  $mean = 0$  and  $std = 1$ .

The CCC loss is used as the loss function. The Adam optimizer with a weight decay of 0.001 is employed. The learning rate (LR) and minimal learning rate (MLR) are set to  $1e-5$  and  $1e-7$ , respectively. The *ReduceLRonPlateau* scheduler with a patience of 5 and factor of 0.1 is employed based on the validation CCC. The maximal epoch number and early stopping counter are set to 100 and 10, respectively. Two groups of layers for the Resnet50 backbone are manually selected for further fine-tuning, which corresponds to the whole layer4 and the last three blocks of layer3.

The training strategy aims to alleviate the out-fitting issue. The Resnet50 backbone is initially fixed except for the output layer. The learning rate is linearly warmed up to  $1e-5$  when  $epoch \leq 10$ . After that, the PyTorch *ReduceLRonPlateau* steps in as the scheduler. The CCC on the validation set determines whether the training improves. If there were no improvements for 5 epochs, the scheduler would reduce the  $LR$  to  $0.1 \cdot LR$ . When  $LR < 1e-7$ , layer4 of the backbone is unfrozen and the scheduler is reset. This process repeats until the last three blocks of layer3 are unfrozen. Finally, when once again  $LR < 1e-7$ , the training is done. Note that at the end of each epoch, the current best model state dictionary (i.e., the one with the greatest validation CCC) is loaded. By doing so, when there is no improvement on validation CCC, the training CCC cannot keep increasing and ends up with over-fitting.

### 3.5. Post-processing

The post-processing consists of CCC-centering and clipping. Given the predictions from 6-fold cross-validation, the CCC-centering aims to yield the weighted prediction based on the inter-class correlation coefficient (ICC) [18]. This technique has been widely used in many emotion recognition challenges [18–21] to obtain the gold-standard labels from multiple raters, by which the bias and inconsistency among individual raters are compensated. The clipping ensures that the inference is truncated within the interval  $[-1, 1]$ , i.e., any values larger or smaller than 1 or  $-1$  are set to 1 or  $-1$ , respectively.

With all the predictions from the 6-fold cross-validation and CCC-centering, we have 5 submissions allowed. They are determined as follows. On the one hand, based on the experience from ABAW2, we find that the influence of the CCC-centering and clipping on the test CCC could be quite trivial ( $\pm 0.001$ ). On the other hand, we also tried other centering methods, such as the estimator weighted evaluator [22] and rater aligned annotation weighting [23]. However, there is almost no visual difference among the three

<sup>2</sup><https://alphacephei.com/vosk/models/vosk-model-en-us-0.22.zip>

<sup>3</sup>[https://docs.nvidia.com/deeplearning/nemo/user-guide/docs/en/main/nlp/punctuation\\_and\\_capitalization.html](https://docs.nvidia.com/deeplearning/nemo/user-guide/docs/en/main/nlp/punctuation_and_capitalization.html)

Table 1. The CCC results from the 6-fold cross-validation on the validation and test sets. Fold 0 is exactly the original data partitioning provided by ABAW3. CCCC denotes CCC-centering. Since 5 submissions are allowed, there are no test results on Fold 2 and 3.

Emotion	Partition	Method	Fold 0	<b>Fold 1</b>	Fold 2	Fold 3	Fold 4	Fold 5	CCCC
Valence	Validation	Baseline	0.310	—	—	—	—	—	—
		Ours	0.450	<b>0.559</b>	0.469	0.531	0.539	0.448	—
	Test	Baseline	0.180	—	—	—	—	—	—
		Ours	0.490	<b>0.520</b>	—	—	0.479	0.511	0.490
Arousal	Validation	Baseline	0.170	—	—	—	—	—	—
		Ours	0.651	<b>0.671</b>	0.564	0.562	0.631	0.618	—
	Test	Baseline	0.170	—	—	—	—	—	—
		Ours	0.584	<b>0.602</b>	—	—	0.580	0.587	0.573

centering methods when we visualize the centered results trial by trial. Finally, our five submissions consist of the predictions from the CCC-centering, Fold0, and the three best folds.

## 4. Result

The validation and test results of our method against the baseline are reported in Table 1. The best results from all the teams are reported in Table 2.

Table 2. The overall test results in CCC. The bold fonts indicate the best results. Citations for several teams are not available from Google Scholar by the time when this report was drafted.

Method	Valence	Arousal	Mean
Situ-RUCAIM3 [24]	<b>0.606</b>	0.596	<b>0.601</b>
Ours	0.520	<b>0.602</b>	0.561
PRL [25]	0.450	0.445	0.448
HSE-NN [26]	0.417	0.454	0.436
AU-NO [27]	0.418	0.407	0.413
LIVIA-2022 [28]	0.374	0.363	0.369
Netease Fuxi Virtual Human [29]	0.300	0.244	0.272
Baseline [2]	0.180	0.170	0.175

## 5. Conclusion

We propose a cross-modal co-attention model for continuous emotion recognition using visual-audio-linguistic information. The model employs three parallel blocks to learn the feature from the three modalities. The learned features are then fused by the cross-modal co-attention block. Experiments are conducted on the Aff-Wild2 database. The achieved CCC on the test set is 0.520 for valence and 0.602 for arousal, which significantly outperforms the baseline method with the corresponding CCC of 0.180 and 0.170 for valence and arousal, respectively.

## References

- [1] G. Sandbach, S. Zafeiriou, M. Pantic, and L. Yin, “Static and dynamic 3d facial expression recognition: A comprehensive survey,” *Image and Vision Computing*, vol. 30, no. 10, pp. 683–697, 2012. 1
- [2] D. Kollias, “Abaw: Valence-arousal estimation, expression recognition, action unit detection & multi-task learning challenges,” *arXiv preprint arXiv:2202.10659*, 2022. 1, 4
- [3] D. Kollias and S. Zafeiriou, “Analysing affective behavior in the second abaw2 competition,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 3652–3660. 1
- [4] D. Kollias, A. Schulc, E. Hajjiev, and S. Zafeiriou, “Analysing affective behavior in the first abaw 2020 competition,” in *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)(FG)*, pp. 794–800. 1
- [5] D. Kollias, V. Sharmanska, and S. Zafeiriou, “Distribution matching for heterogeneous multi-task learning: a large-scale face study,” *arXiv preprint arXiv:2105.03790*, 2021. 1
- [6] D. Kollias and S. Zafeiriou, “Affect analysis in-the-wild: Valence-arousal, expressions, action units and a unified framework,” *arXiv preprint arXiv:2103.15792*, 2021. 1
- [7] —, “Expression, affect, action unit recognition: Aff-wild2, multi-task learning and arcface,” *arXiv preprint arXiv:1910.04855*, 2019. 1
- [8] D. Kollias, V. Sharmanska, and S. Zafeiriou, “Face behavior a la carte: Expressions, affect and action units in a single network,” *arXiv preprint arXiv:1910.11111*, 2019. 1

- [9] D. Kollias, P. Tzirakis, M. A. Nicolaou, A. Papaioannou, G. Zhao, B. Schuller, I. Kotsia, and S. Zafeiriou, “Deep affect prediction in-the-wild: Aff-wild database and challenge, deep architectures, and beyond,” *International Journal of Computer Vision*, pp. 1–23, 2019. [1](#)
- [10] S. Zafeiriou, D. Kollias, M. A. Nicolaou, A. Papaioannou, G. Zhao, and I. Kotsia, “Aff-wild: Valence and arousal in-the-wild challenge,” in *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*. IEEE, 2017, pp. 1980–1987. [1](#)
- [11] S. Zhang, Y. Ding, Z. Wei, and C. Guan, “Continuous emotion recognition with audio-visual leader-follower attentive fusion,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 3567–3574. [1](#)
- [12] S. Hershey, S. Chaudhuri, D. P. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold *et al.*, “Cnn architectures for large-scale audio classification,” in *2017 IEEE international conference on acoustics, speech and signal processing (icassp)*. IEEE, 2017, pp. 131–135. [1](#)
- [13] S. Bai, J. Z. Kolter, and V. Koltun, “An empirical evaluation of generic convolutional and recurrent networks for sequence modeling,” *arXiv preprint arXiv:1803.01271*, 2018. [1](#)
- [14] E. Barsoum, C. Zhang, C. C. Ferrer, and Z. Zhang, “Training deep networks for facial expression recognition with crowd-sourced label distribution,” in *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, 2016, pp. 279–283. [1](#)
- [15] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, “Language models are unsupervised multitask learners,” *OpenAI blog*, vol. 1, no. 8, p. 9, 2019. [1](#)
- [16] R. Cowie, E. Douglas-Cowie, S. Savvidou\*, E. McMahon, M. Sawey, and M. Schröder, “‘feeltrace’: An instrument for recording perceived emotion in real time,” in *ISCA tutorial and research workshop (ITRW) on speech and emotion*, 2000. [2](#)
- [17] L. Sun, Z. Lian, J. Tao, B. Liu, and M. Niu, “Multi-modal continuous dimensional emotion recognition using recurrent neural network and self-attention mechanism,” in *Proceedings of the 1st International on Multimodal Sentiment Analysis in Real-life Media Challenge and Workshop*, 2020, pp. 27–34. [3](#)
- [18] F. Ringeval, B. Schuller, M. Valstar, N. Cummins, R. Cowie, L. Tavabi, M. Schmitt, S. Alisamir, S. Amiriparian, E.-M. Messner *et al.*, “Avec 2019 workshop and challenge: state-of-mind, detecting depression with ai, and cross-cultural affect recognition,” in *Proceedings of the 9th International on Audio/Visual Emotion Challenge and Workshop*, 2019, pp. 3–12. [3](#)
- [19] M. Valstar, J. Gratch, B. Schuller, F. Ringeval, D. Lalanne, M. Torres Torres, S. Scherer, G. Stratou, R. Cowie, and M. Pantic, “Avec 2016: Depression, mood, and emotion recognition workshop and challenge,” in *Proceedings of the 6th international workshop on audio/visual emotion challenge*, 2016, pp. 3–10. [3](#)
- [20] F. Ringeval, B. Schuller, M. Valstar, J. Gratch, R. Cowie, S. Scherer, S. Mozgai, N. Cummins, M. Schmitt, and M. Pantic, “Avec 2017: Real-life depression, and affect recognition workshop and challenge,” in *Proceedings of the 7th Annual Workshop on Audio/Visual Emotion Challenge*, 2017, pp. 3–9. [3](#)
- [21] F. Ringeval, B. Schuller, M. Valstar, R. Cowie, H. Kaya, M. Schmitt, S. Amiriparian, N. Cummins, D. Lalanne, A. Michaud *et al.*, “Avec 2018 workshop and challenge: Bipolar disorder and cross-cultural affect recognition,” in *Proceedings of the 2018 on audio/visual emotion challenge and workshop*, 2018, pp. 3–13. [3](#)
- [22] B. W. Schuller, *Intelligent audio analysis*. Springer, 2013. [3](#)
- [23] L. Stappen, L. Schumann, B. Sertolli, A. Baird, B. Weigell, E. Cambria, and B. W. Schuller, “Muse-toolbox: The multimodal sentiment analysis continuous annotation fusion and discrete class transformation toolbox,” in *Proceedings of the 2nd on Multimodal Sentiment Analysis Challenge*, 2021, pp. 75–82. [3](#)
- [24] L. Meng, Y. Liu, X. Liu, Z. Huang, W. Jiang, T. Zhang, Y. Deng, R. Li, Y. Wu, J. Zhao *et al.*, “Multi-modal emotion estimation for in-the-wild videos,” *arXiv preprint arXiv:2203.13032*, 2022. [4](#)
- [25] H.-H. Nguyen, V.-T. Huynh, and S.-H. Kim, “An ensemble approach for facial expression analysis in video,” *arXiv preprint arXiv:2203.12891*, 2022. [4](#)
- [26] A. V. Savchenko, “Frame-level prediction of facial expressions, valence, arousal and action units for mobile devices,” *arXiv preprint arXiv:2203.13436*, 2022. [4](#)
- [27] V. Karas, M. K. Tellamekala, A. Mallol-Ragolta, M. Valstar, and B. W. Schuller, “Continuous-time

audiovisual fusion with recurrence vs. attention for in-the-wild affect recognition,” *arXiv preprint arXiv:2203.13285*, 2022. 4

[28] G. P. Rajasekar, W. C. de Melo, N. Ullah, H. Aslam, O. Zeeshan, T. Denorme, M. Pedersoli, A. Koerich, P. Cardinal, and E. Granger, “A joint cross-attention model for audio-visual fusion in dimensional emotion recognition,” *arXiv preprint arXiv:2203.14779*, 2022. 4

[29] W. Zhang, Z. Zhang, F. Qiu, S. Wang, B. Ma, H. Zeng, R. An, and Y. Ding, “Transformer-based multimodal information fusion for facial expression analysis,” *arXiv preprint arXiv:2203.12367*, 2022. 4