

Multimodal Continuous Emotion Recognition: A Technical Report for ABAW5

Su Zhang¹, Ziyuan Zhao^{1,2}, Cuntai Guan^{1,*}

¹Nanyang Technological University

²Institute of Infocomm Research (I²R), A*STAR, Singapore

sorazcn@gmail.com, S210088@e.ntu.edu.sg, ctguan@ntu.edu.sg

Abstract

We used two multimodal models for continuous valence-arousal recognition using visual, audio, and linguistic information. The first model is the same as we used in ABAW2 and ABAW3, which employs the leader-follower attention. The second model has the same architecture for spatial and temporal encoding. As for the fusion block, it employs a compact and straightforward channel attention, borrowed from the End2You toolkit. Unlike our previous attempts that use Vggish feature directly as the audio feature, this time we feed the pre-trained VGG model using logmel-spectrogram and finetune it during the training. To make full use of the data and alleviate over-fitting, cross-validation is carried out. 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 [18], 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 . The term "continuous" is endowed with two distinctive features. Spatially, it seeks to represent the emotional state as a point with continuous numerical values in the multi-dimensional space of the dimensional theory, as opposed to assigning categorical labels. Temporally, it incessantly forecasts the emotional state for a predetermined duration, forming a continuous record of the subject's emotional trajectory over a specified period.

The task of Continuous Emotion Recognition (CER) is challenging, mainly due to the following reasons. Firstly,

the complexity of emotions over an extended period can be seen as a composition of multiple facial muscular movements, which can be defined as one-actions [5] using the Facial Action Coding System (FACS) [19]. However, despite its atomic nature, human emotions, irrespective of the modality, typically exhibit large variations in their intensity, order, and duration, and take longer to unfold. Consequently, there is a need for models that can learn long-range temporal dependencies to address this issue. Secondly, emotions are highly subjective and vary significantly depending on individual experiences, making their perception susceptible to personal bias. For instance, children who have been subjected to physical abuse can detect signs of anger much faster than their counterparts [15]. Therefore, the data collected from subjects and the ground truth provided by annotators are prone to individual biases. One promising technique to mitigate this issue is multimodality, which involves the integration of visual, audio, and physiological data to develop reliable CER models.

In order to address the challenges posed by temporal dynamic learning and cross-subject generality, we adopt two strategies: large window resampling and multimodal fusion. Large window resampling involves selecting a window size that determines the amount of contextual information the model considers to make a prediction for each time step. On the other hand, multimodal fusion is a widely accepted technique that leverages the complementary nature of multiple data modalities to improve prediction accuracy. This is because multimodal data can disambiguate conflicting emotional cues, such as a crying face with joyful vocal expressions, or a neutral face with a harsh intonation. Additionally, multimodal data can enhance recognition robustness in real-world scenarios characterized by uncontrolled environments, subject diversity, various illumination conditions, spontaneous behaviors, background noise, and unclear speech. For instance, when an actor's face is not visible, their voice or body gesture can serve as an emotional cue to support ongoing learning.

This report details our methodology for the valence-arousal estimation challenge from the fifth affective behav-

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

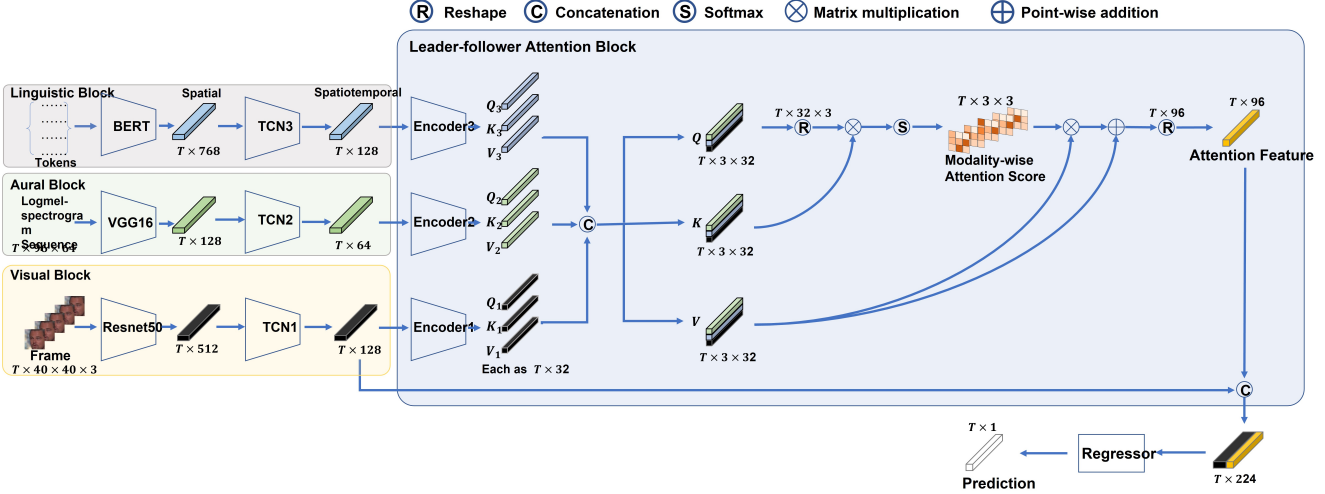


Figure 1. The architecture of our proposed model. The model consists of four components, i.e., the visual, audio, linguistic, and co-attention blocks. The visual block has a cascaded 2DCNN-TCN structure, and the audio and linguistic blocks each contain a TCN. The three branches yield three independent 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 Q_i , a key K_i , and a value V_i . They are then regrouped and concatenated to form the cross-modal counterparts.

ior analysis in-the-wild (ABAW5) workshop [6–9, 11–14, 25]. ABAW5 aims for the affective behavior analysis in-the-wild. The valence-arousal estimation sub-challenge bases on the extended version of the Aff-Wild2 database, in which the valence and arousal are annotated continuously for each video frame. Our work is an extension of the last year’s attempt [26] on ABAW5 [14]. The extension are twofold. First, a pre-trained audio VGGnet is employed to serve as the audio backbone for the vggish [4] extraction on-the-go. Thus, the expressiveness of deep feature on visual and audio modalities can be further improved through finetuning. Second, a compact and straightforward fusion block is borrowed from the End2You toolkit, in which the channel attention is employed to fuse the multimodal information. The results of using leader-follower attention and channel attention are reported.

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. Section 5 concludes the work.

2. Model Architecture

There are two architectures for our model, the leader-follower attention network (LFAN) and channel attention network (CAN). Both of them consists of independent branches to extract the spatiotemporal encoding for

each modality, followed by the fusion block using leader-follower attention and channel attention, respectively. For each branch, a CNN backbone is employed to extract the spatial deep feature on-the-go for video frames and logmel-spectrogram, followed by a temporal convolutional network (TCN) [1] to further learn the spatiotemporal encoding. For the linguistic branch, the pre-extracted bert features [17] are taken as the input to feed the TCN. For branches fed by low-level features, no backbone is employed as well. The illustration of LFAN using video frames, logmel-spectrograms, and speech tokens are shown in Fig 1. As for the details of the channel attention-based fusion block, please refer to the End2You paper [22] and the code at <https://github.com/end2you/end2you>.

3. Implementation Details

3.1. Database

The ABAW5 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 [2]. 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 $48 \times 48 \times 3$ as a compromise to limited computational power. 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 logmel-spectrogram are then extracted using the preprocessing code from the Vggish repository¹. The only change is that we specified the hop length to be $1/fps$ 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 toolkit², from which the recognized words and the word-level timestamp are obtained. The recognized words are then fed to a pretrained punctuation model from Hugging Face³. 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 [21]. 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 steps 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.

The audio preprocessing firstly converts all the videos to mono with a $16K$ sampling rate in wav format. After which, the logmel-spectrogram is extracted following the Vggish preprocessing code⁴. Note that the hop length for these features are set to be $1/fps$ 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 toolkit⁵, 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 deepmultilingualpunctuation toolkit⁶. 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 [21]. 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 steps 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 and annotations have the same length, the feature matrices are either repeatedly padded (using the last feature points) or trimmed, depending on whether the feature length is shorter and longer than the trial length, respectively.

3.3. Data Expansion

The AffWild2 database employed by ABAW5 contains 360, 72, and 162 trials in the training, validation, and testing sets, respectively. To make full use of the available data and alleviate over-fitting, 6-fold cross-validation is employed. By evenly splitting the training set into 5 folds, we have 6×72 trials in total. Note that the 0-th fold is exactly the original data partitioning. And there is no subject overlap across different folds. Moreover, during training and validation, the resampling window has a 33% overlap, resulting in 33% more data.

3.4. Training

The batch size is 12. 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 bert features, 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 maximal epoch number and early stopping counter are set to 100 and 20, respectively. The learning rate (LR) and min-

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

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

³<https://huggingface.co/felflare/bert-restore-punctuation>

⁴<https://github.com/harritaylor/torchvggish>

⁵<https://alphacephei.com/vosk/models/vosk-model-en-us-0.22.zip>

⁶<https://pypi.org/project/deepmultilingualpunctuation/>

Table 1. The CCC results from the 6-fold cross-validation on the validation set. Fold 0 is exactly the original data partitioning provided by ABAW5.

Emotion	Method	Fold 0	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5
Valence	LFAN	0.441	—	—	—	—	—
	CAN	0.423	0.612	0.484	0.529	0.543	0.488
Arousal	LFAN	0.645	—	—	—	—	—
	CAN	0.670	0.680	0.592	0.618	0.655	0.610

imal learning rate (MLR) are set to $1e - 5$ and $1e - 8$, respectively.

The warmup scheme with *ReduceLRonPlateau* scheduler with a patience of 5 and factor of 0.1 is employed based on the validation CCC. Three groups of layers for the visual Resnet50 backbone and audio VGG backbone are manually selected for further fine-tuning. When epoch= 0, the first group is unfrozen. The learning rate is then linearly warmed up to $1e - 5$ within an epoch. The repetitive warm-up is carried out until epoch= 5. After which the *ReduceLRonPlateau* takes over to update the learning rate. It gradually drops the learning rate in a factor of 0.1 should no higher validation CCC appears for a consecutive 5 epochs. When the learning rate is already $1e - 8$ and no higher validation CCC appears for the latest 5 epochs, the second group is unfrozen, the learning rate is also reset to $1e - 5$ before undergoing the warm-up scheme and *ReduceLRonPlateau* again. The procedure is repeated until no groups are available to unfrozen. 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.

Note that unlike our previous attempts on ABAW2 and ABAW3, no post-processing is employed. Given the results from 6 folds, we simply choose the ones with the highest validation CCC to submit. Also note that the results for any fold are selected from 15 instances using different seeds.

4. Result and Discussion

The validation results of our methods against the baseline are reported in Table 1. Note that for one emotion type, the reported result for each fold is the best instance selected from 15 instances using different seeds. Finally, the best results of Fold 0 from CAN and LFAN, together with other 3 best folds (measured in the average CCC) from CAN, are submitted to ABAW5.

In this study, two multimodal models are compared, namely CAN and LFAN, with the former utilizing only video frames and logmel-spectrograms while the latter incorporates an additional Bert feature. Validation results indicate that the CAN model with two modalities outperforms

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
SituTech	0.619	0.663	0.641
Netease Fuxi Virtual Human [27]	0.649	0.626	0.637
Ours	0.553	0.630	0.591
CtyunAI [29]	0.501	0.633	0.567
HFUT-MAC [28]	0.523	0.545	0.534
HSE-NN-SberAI [20]	0.482	0.528	0.505
ACCC	0.462	0.506	0.482
PRL [23]	0.504	0.428	0.466
SCLAB CNU [16]	0.458	0.470	0.464
USTC-AC [24]	0.325	0.232	0.278
Baseline [10]	0.211	0.191	0.201

the LFAN model with three modalities, achieving a higher CCC score in Fold 0. It is worth noting that attempts to include the Bert feature as a third modality in the CAN model did not result in improvement. This suggests that the current linguistic information encoded by the pre-trained Bert model does not provide useful information to the model, or at least not in its current form. The issue lies in the fact that the length of tokens is not equivalent to the length of the sampling window, rendering fine-tuning impossible during training. Future research is needed to explore how linguistic information can be effectively incorporated into the fusion process. Furthermore, the study highlights the efficacy of fine-tuning. Both Resnet50 and VGG16 models, trained using large public datasets and fed by video frames and logmel-spectrograms, show improvements in performance when the last few layers are fine-tuned during training. This improvement can be as high as 10% or more, compared to models that use pre-extracted CNN and Vggish features, albeit at a cost of 300% – 500% longer training time.

The best results from all the teams are reported in Table 2.

5. Conclusion

In our study, we aimed to develop effective models for continuous valence-arousal recognition, utilizing multimodal information from visual, audio, and linguistic sources. To overcome challenges posed by temporal dynamic learning and cross-subject generality, we employed two strategies: large window resampling and multimodal fusion. Specifically, we designed the CAN and LFAN models, with a resampling window length and hop length of 300 and 200, respectively, and multimodal fusion block to effectively integrate information from different modalities. Moving forward, our research will explore the incorporation of transformer-based networks for improved temporal dynamics learning.

References

- [1] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018. 2
- [2] Roddy Cowie, Ellen Douglas-Cowie, Susie Savvidou*, Edelle McMahon, Martin Sawey, and Marc Schröder. 'feel-trace': An instrument for recording perceived emotion in real time. In *ISCA tutorial and research workshop (ITRW) on speech and emotion*, 2000. 2
- [3] Florian Eyben, Martin Wöllmer, and Björn Schuller. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM international conference on Multimedia*, pages 1459–1462, 2010.
- [4] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al. Cnn architectures for large-scale audio classification. In *2017 IEEE international conference on acoustics, speech and signal processing (icassp)*, pages 131–135. IEEE, 2017. 2
- [5] Noureldien Hussein, Efstratios Gavves, and Arnold WM Smeulders. Timeception for complex action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 254–263, 2019. 1
- [6] Dimitrios Kollias. Abaw: Valence-arousal estimation, expression recognition, action unit detection & multi-task learning challenges. *arXiv preprint arXiv:2202.10659*, 2022. 2
- [7] D Kollias, A Schulc, E Hajiyev, 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)*, pages 794–800. 2
- [8] Dimitrios Kollias, Viktoriia Sharmanska, and Stefanos Zafeiriou. Face behavior a la carte: Expressions, affect and action units in a single network. *arXiv preprint arXiv:1910.11111*, 2019. 2
- [9] Dimitrios Kollias, Viktoriia Sharmanska, and Stefanos Zafeiriou. Distribution matching for heterogeneous multi-task learning: a large-scale face study. *arXiv preprint arXiv:2105.03790*, 2021. 2
- [10] Dimitrios Kollias, Panagiotis Tzirakis, Alice Baird, Alan Cowen, and Stefanos Zafeiriou. Abaw: Valence-arousal estimation, expression recognition, action unit detection & emotional reaction intensity estimation challenges, 2023. 4
- [11] Dimitrios Kollias, Panagiotis Tzirakis, Mihalis A Nicolaou, Athanasios Papaioannou, Guoying Zhao, Björn Schuller, Irene Kotsia, and Stefanos Zafeiriou. Deep affect prediction in-the-wild: Aff-wild database and challenge, deep architectures, and beyond. *International Journal of Computer Vision*, pages 1–23, 2019. 2
- [12] Dimitrios Kollias and Stefanos Zafeiriou. Expression, affect, action unit recognition: Aff-wild2, multi-task learning and arcfac. *arXiv preprint arXiv:1910.04855*, 2019. 2
- [13] Dimitrios Kollias and Stefanos Zafeiriou. Affect analysis in-the-wild: Valence-arousal, expressions, action units and a unified framework. *arXiv preprint arXiv:2103.15792*, 2021. 2
- [14] Dimitrios Kollias and Stefanos Zafeiriou. Analysing affective behavior in the second abaw2 competition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3652–3660, 2021. 2
- [15] David G Myers. *Psychology*, 2004. 1
- [16] Dang-Khanh Nguyen, Ngoc-Huynh Ho, Sudarshan Pant, and Hyung-Jeong Yang. A transformer-based approach to video frame-level prediction in affective behaviour analysis in-the-wild, 2023. 4
- [17] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 2
- [18] Georgia Sandbach, Stefanos Zafeiriou, Maja Pantic, and Lijun Yin. Static and dynamic 3d facial expression recognition: A comprehensive survey. *Image and Vision Computing*, 30(10):683–697, 2012. 1
- [19] Nishant Sankaran, Deen Dayal Mohan, Nagashri N Lakshminarayana, Srirangaraj Setlur, and Venu Govindaraju. Domain adaptive representation learning for facial action unit recognition. *Pattern Recognition*, 102:107127, 2020. 1
- [20] Andrey V. Savchenko. Emotiefnet facial features in uni-task emotion recognition in video at abaw-5 competition, 2023. 4
- [21] Licai Sun, Zheng Lian, Jianhua Tao, Bin Liu, and Mingyue 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*, pages 27–34, 2020. 3
- [22] Panagiotis Tzirakis, Stefanos Zafeiriou, and Bjorn W Schuller. End2you—the imperial toolkit for multimodal profiling by end-to-end learning. *arXiv preprint arXiv:1802.01115*, 2018. 2
- [23] Tu Vu, Van Thong Huynh, and Soo Hyung Kim. Vision transformer for action units detection, 2023. 4
- [24] Shangfei Wang, Yanan Chang, Yi Wu, Xiangyu Miao, Jiaqiang Wu, Zhouan Zhu, Jiahe Wang, and Yufei Xiao. Facial affective behavior analysis method for 5th abaw competition, 2023. 4

- [25] Stefanos Zafeiriou, Dimitrios Kollias, Mihalis A Nicolaou, Athanasios Papaioannou, Guoying Zhao, and Irene Kotsia. Aff-wild: Valence and arousal ‘in-the-wild’ challenge. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*, pages 1980–1987. IEEE, 2017. 2
- [26] Su Zhang, Ruyi An, Yi Ding, and Cuntai Guan. Continuous emotion recognition using visual-audio-linguistic information: A technical report for abaw3. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2376–2381, 2022. 2
- [27] Wei Zhang, Bowen Ma, Feng Qiu, and Yu Ding. Multi-modal facial affective analysis based on masked autoencoder, 2023. 4
- [28] Ziyang Zhang, Liuwei An, Zishun Cui, Ao xu, Tengting Dong, Yueqi Jiang, Jingyi Shi, Xin Liu, Xiao Sun, and Meng Wang. Facial affect recognition based on transformer encoder and audiovisual fusion for the abaw5 challenge, 2023. 4
- [29] Weiwei Zhou, Jiada Lu, Zhaolong Xiong, and Weifeng Wang. Continuous emotion recognition based on ten and transformer, 2023. 4