

Multi-Task Transformer with uncertainty modelling for Face Based Affective Computing

Gauthier Tallec
ISIR

Jules Bonnard
ISIR

Arnaud Dapogny
Datakalab

Kevin Bailly
ISIR/Datakalab

Abstract

Face based affective computing consists in detecting emotions from face images. It is useful to unlock better automatic comprehension of human behaviours and could pave the way toward improved human-machines interactions. However it comes with the challenging task of designing a computational representation of emotions. So far, emotions have been represented either continuously in the 2D Valence/Arousal (VA) space or in a discrete manner with Ekman’s 7 basic emotions (FER). Alternatively, Ekman’s Facial Action Unit (AU) system have also been used to characterize emotions using a codebook of unitary muscular activations. ABAW3 and ABAW4 Multi-Task Challenges are the first work to provide a large scale database annotated with those three types of labels. In this paper we present a transformer based multi-task method for jointly learning to predict valence arousal, action units and basic emotions. From an architectural standpoint our method uses a task-wise token approach to efficiently model the similarities between the tasks. From a learning point of view we use an uncertainty weighted loss for modelling the difference of stochasticity between the three tasks annotations.

1. Introduction

The Affective Behavior Analysis In-The-Wild (ABAW) competitions are a serie of machine learning based affective computing challenges based on the Aff-Wild2 database [12, 17]. In the past few years, the competition objectives have focused on learning to predict different face expression representation [9, 15] either jointly [8, 10, 11, 14] or in a separate manner [4]. This year is no exception, as ABAW4 proposes (a) A challenge for Multi-Task learning that consists in predicting Valence-Arousal, Facial Expression, and Action Units [5] (b) A challenge for learning to predict the 7 basic emotions from synthetic data [3, 6, 7, 13].

We participate in the multi-task learning challenge because this is one of the first work to provide a large scale

database (170m frames) annotated in all three emotion representations (AU, VA and FER).

One of the main challenge when learning to predict emotion representations is that the provided data are often noisy due to the intrinsic subjectivity of the annotation task. It is all the more a problem when trying to jointly predict several emotion representations as those representations may display different level of noise due to their different level of subjectivity. For example Valence Arousal annotations are slightly noisier than Action Units due to the continuous nature of the annotations. To tackle that challenge we leverage a classic multi-task transformer architecture that we supervise with an uncertainty weighted loss [2] that takes into account the different stochasticity levels of each task.

2. Related Work

The lack of annotated datasets in different emotional representations make multi-task face analysis models rarer in the literature. Nonetheless, the recent ABAW challenges have increased the interest in the multi-task approach to predict the three emotion descriptors.

For the ABAW3 challenge, Didan Deng [1] introduced a ”Sign-and-Message” model which uses an Emotion Transformer to predict Action Units and metric learning to predict facial expressions, valence and arousal.

Jeong *et al.*, the runner-ups of the ABAW3 Multi-Task challenge jointly leveraged the video data with an LSTM and the audio data with a SoundNet. These latent representations are concatenated and fed to different task classifiers and regressors. A task discriminator is used to improve the individual task predictions.

However, the ABAW4 challenge doesn’t provide any audio data and the available video data and the important number of missing frames in each video may prevent the use from of temporal methods.

3. Methodology

In this section we build up our transformer architecture based on the background material provided in [16]. Fur-

thermore we provide ourselves with a multi-task dataset $\mathcal{D} = \{\mathbf{x}_i, \mathbf{v}_i, \mathbf{a}_i, \mathbf{e}_i\}_{i=1}^N$ where $\mathbf{x} \in \mathbb{R}^{H \times W \times C}$ is the input image, $\mathbf{v} \in \mathbb{R}^2$ is the valence arousal ground truth, $\mathbf{a} \in \{0, 1\}^{12}$ is the AU ground truth and $\mathbf{e} \in \{0, 1\}^8$ is the basic emotion where the 8th label features the additional 'Other' class that is present in ABAW4.

3.1. General Architecture

ViT-like architecture are known to be sensitive to overfitting and to therefore require a lot of reliable data to be able to overcome the performance of basic convolutional networks.

For that purpose we adopt an hybrid architecture in which we first feed the input image to a convolutional encoder which consist in a convolutional encoder in which the final pooling layer is replaced by a 1D convolution with output dimension d .

Each of the $H' \times W'$ output patch representations are then added an usual cos-sin based encoding to account for its position in the image. Those patches are then passed through N_x classic self attention layers resulting in a patch based image representation that we denote $\mathbf{p} \in \mathbb{R}^{H' \times W' \times d}$ the resulting image representation.

To predict each of the three tasks we provide ourselves with three learned task tokens. Those tokens are then fed to a sequential arrangement of N_t attention modules. Each cross attention module is composed of two main block:

1. The first block consist in self-attention between tokens so that each task can learn to exploit the information from the other tasks.
2. The second block is a cross-attention step in which each token learns which part of the \mathbf{p} is most useful to predict its associated task.

At the end of the N_t modules, each refined token is fed to a sequence of N_d task specific dense layers where all activations are relu except for the last one which is linear.

The final output of this process is a prediction $\hat{\mathbf{v}}$ for valence arousal, and logits \mathbf{u}_a and \mathbf{u}_e for action units and emotion respectively. The predictions for actions units $\hat{\mathbf{a}}$ and emotion $\hat{\mathbf{e}}$ are computed using different temperature T_a , T_e to account for the different level of stochasticity of those two tasks:

$$\hat{\mathbf{a}} = \sigma\left(\frac{\mathbf{u}_a}{T_a}\right), \hat{\mathbf{e}} = S\left(\frac{\mathbf{u}_e}{T_e}\right), \quad (1)$$

where σ and S denotes sigmoid and softmax function respectively.

3.2. Uncertainty Based Supervision

To account for task stochasticity we assigned temperatures T_a and T_e to the action unit and emotion prediction task respectively. Similarly we assign a variance matrix to arousal

and valence prediction. To simplify we work under the assumption that this variance matrix is diagonal and that valence and arousal have the same intrinsic variance σ_v^2 . The final Maximum likelihood based supervision loss is then :

$$\mathcal{L} = \mathcal{L}_v + \mathcal{L}_a + \mathcal{L}_e \quad (2)$$

where :

$$\mathcal{L}_v = \frac{1}{2\sigma_v^2} \text{MSE}(\hat{\mathbf{v}}, \mathbf{v}), \quad (3)$$

$$\mathcal{L}_a = \frac{1}{2T_a} \text{BCE}(\mathbf{u}_a, \mathbf{a}), \quad (4)$$

$$\mathcal{L}_e = \frac{1}{T_e} \text{CCE}(\mathbf{u}_e, \mathbf{e}), \quad (5)$$

and MSE, BCE and CCE denotes mean squared error, binary crossentropy and categorical crossentropy respectively. In practice, as all three labels are not present for each examples we mask the losses for the missing labels and reweight the total loss with the number of present labels so that each example get the same contribution in the loss.

4. Experiments

4.1. Implementation Details

For all our experiments we use a Resnet50 pretrained on VGGFACE2 as the convolutional backbone. For each resized image of size $224 \times 224 \times 3$, it outputs a 7×7 patch based representation that we project in $d = 768$ dimensions using a 1D convolution. We then apply $N_x = 2$ self attention layers with 12 attention heads to get the final image representation \mathbf{p} .

As far as the task-token based part is concerned it consists in $N_t = 2$ cross attention modules with 12 attention heads followed by $N_d = 4$ taskwise dense.

The variance σ_v^2 of valence arousal is set to 1. The temperature of action units and emotion are respectively set to $T_a = 1$ and $T_e = 5$.

4.2. Data Augmentation

ABAW4 dataset features in the wild data, so that there is a lot of variability in images and potentially between the train, the valid and the test set. To reduce the impact of those variabilities we use classical geometrical data augmentation : random rotations, translations and zooms. Furthermore we tried using multi-task MixUp which consist in feeding random convex combinations of input images and supervising each task with the same random convex combination of each task ground truth.

4.3. Temporal Smoothing

The provided test set consists in images sampled from videos. To handle the dynamic of labels we simply compute the final predictions as temporal means of predictions

| ID | Mixup | AU F1 | Emotion F1 | VA CCC | ABAW4 |
|----|-------|-------|------------|--------|-------|
| 1 | × | 0.25 | 0.48 | 0.39 | 1.12 |
| 2 | ✓ | 0.27 | 0.47 | 0.38 | 1.11 |

Table 1. Summary of validation scores for submissions to the ABAW4 Challenge

in windows of size $S = 30$. For valence/arousal, the predictions are directly smoothed. For AU and Emotions, the logits are smoothed and the appropriate activation function is applied on the smoothing output. Missing images in the videos are simply ignored in the mean computation.

Table 4.3 summarize the validation scores of our two submissions. Our third submission is simply the mean of the two first submission.

5. Conclusion

In this work we made use of a multi-task transformer for jointly predicting arousal/valence, basic emotions and action units. We further refined our method by modelling the different levels of stochasticity for each task. Our method finished at the 8th position of the ABAW4 Challenge.

6. Acknowledgements

This work was granted access to the HPC resources of IDRIS under the allocation 2021-AD011013183 made by GENCI.

References

- [1] Didan Deng. Multiple emotion descriptors estimation at the abaw3 challenge, 2022.
- [2] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018.
- [3] Dimitrios Kollias. Abaw: Learning from synthetic data & multi-task learning challenges. *arXiv preprint arXiv:2207.01138*, 2022.
- [4] Dimitrios Kollias. Abaw: Valence-arousal estimation, expression recognition, action unit detection & multi-task learning challenges. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2328–2336, 2022.
- [5] Dimitrios Kollias. Abaw: Valence-arousal estimation, expression recognition, action unit detection & multi-task learning challenges. *arXiv preprint arXiv:2202.10659*, 2022.
- [6] Dimitrios Kollias, Shiyang Cheng, Maja Pantic, and Stefanos Zafeiriou. Photorealistic facial synthesis in the dimensional affect space. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018.
- [7] Dimitrios Kollias, Shiyang Cheng, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. Deep neural network augmentation: Generating faces for affect analysis. *International Journal of Computer Vision*, 128(5):1455–1484, 2020.
- [8] Dimitrios Kollias, Mihalios A Nicolaou, Irene Kotsia, Guoying Zhao, and Stefanos Zafeiriou. Recognition of affect in the wild using deep neural networks. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*, pages 1972–1979. IEEE, 2017.
- [9] D Kollias, A Schulc, E Hajiyeve, 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.
- [10] 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.
- [11] 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.
- [12] Dimitrios Kollias, Panagiotis Tzirakis, Mihalios 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.
- [13] Dimitrios Kollias and Stefanos Zafeiriou. Va-stargan: Continuous affect generation. In *International Conference on Advanced Concepts for Intelligent Vision Systems*, pages 227–238. Springer, 2020.
- [14] 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.
- [15] 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.
- [16] Gauthier Tallec, Edouard Yvinec, Arnaud Dapogny, and Kevin Bailly. Multi-label transformer for action unit detection. *arXiv preprint arXiv:2203.12531*, 2022.
- [17] Stefanos Zafeiriou, Dimitrios Kollias, Mihalios 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.