

COLD Fusion: Calibrated and Ordinal Latent Distribution Fusion for Uncertainty-Aware Multimodal Emotion Recognition

Mani Kumar Tellamekala¹, Shahin Amiriparian¹, *Member, IEEE*, Björn W. Schuller², *Fellow, IEEE*, Elisabeth André¹, Timo Giesbrecht¹, and Michel Valstar¹

Abstract—Automatically recognising apparent emotions from face and voice is hard, in part because of various sources of uncertainty, including in the input data and the labels used in a machine learning framework. This paper introduces an uncertainty-aware multimodal fusion approach that quantifies modality-wise aleatoric or data uncertainty towards emotion prediction. We propose a novel fusion framework, in which latent distributions over unimodal temporal context are learned by constraining their variance. These variance constraints, Calibration and Ordinal Ranking, are designed such that the variance estimated for a modality can represent how informative the temporal context of that modality is w.r.t. emotion recognition. When well-calibrated, modality-wise uncertainty scores indicate how much their corresponding predictions are likely to differ from the ground truth labels. Well-ranked uncertainty scores allow the ordinal ranking of different frames across different modalities. To jointly impose both these constraints, we propose a softmax distributional matching loss. Our evaluation on AVEC 2019 CES, CMU-MOSEI, and IEMOCAP datasets shows that the proposed multimodal fusion method not only improves the generalisation performance of emotion recognition models and their predictive uncertainty estimates, but also makes the models robust to novel noise patterns encountered at test time.

Index Terms—Dimensional affect recognition, multimodal fusion, uncertainty modeling, categorical emotion recognition.

Manuscript received 10 January 2023; revised 17 August 2023; accepted 12 October 2023. Date of publication 18 October 2023; date of current version 8 January 2024. The work of Mani Tellamekala was supported in part by Engineering and Physical Science Research Council project under Grant 2159382, and in part by Unilever U.K. Ltd., and the work of Michel Valstar’s work was supported by Nottingham Biomedical Research Centre. Recommended for acceptance by V. Patel. (*Corresponding author: Mani Kumar Tellamekala.*)

Mani Kumar Tellamekala and Michel Valstar are with the Computer Vision Lab, School of Computer Science, University of Nottingham, NG7 2RD Nottingham, U.K. (e-mail: mani.kumar.t@outlook.com; michel.valstar@nottingham.ac.uk).

Shahin Amiriparian is with the Chair of Embedded Intelligence for Health Care & Wellbeing, University of Augsburg, 86159 Augsburg, Germany (e-mail: shahin.amiriparian@informatik.uni-augsburg.de).

Björn W. Schuller is with the Chair of Embedded Intelligence for Health Care & Wellbeing, University of Augsburg, 86159 Augsburg, Germany, and also with GLAM – the Group on Language, Audio, & Music, Imperial College London, SW7 2BX London, U.K. (e-mail: schuller@informatik.uni-augsburg.de).

Elisabeth André is with the Chair for Human-Centered Artificial Intelligence, University of Augsburg, 86159 Augsburg, Germany (e-mail: andre@informatik.uni-augsburg.de).

Timo Giesbrecht is with Unilever R&D Port Sunlight, CH62 4UY Wirral, U.K. (e-mail: timo.giesbrecht@unilever.com).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TPAMI.2023.3325770>, provided by the authors.

Digital Object Identifier 10.1109/TPAMI.2023.3325770

I. INTRODUCTION

LEARNING to fuse task-specific information from multiple modalities is a fundamental problem in Machine Learning. At its core, this problem entails estimating how informative each modality is towards predicting the labels of a target task. For example, consider the task of automatically recognising emotional expressions from a video in which a person is talking with a face mask covering. In such a scenario, for effectively fusing information from the audio and visual modalities, the model must be aware of how informative the facial and vocal streams are w.r.t the target task separately. Thus, modality-wise uncertainty-aware fusion is a natural approach to multimodal learning.

In this work, we formulate an uncertainty-aware fusion method for the task of apparent emotion recognition from multimodal inputs. The proposed multimodal fusion framework is based on probabilistic modelling of unimodal temporal context related to emotional expressions. This probabilistic temporal modelling approach aims to capture the richness of the temporal context in terms of emotional expressions present in a given modality, and use that information in deciding the degree of importance of each modality towards recognising apparent emotions.

In the proposed method, we first estimate the uncertainty of unimodal temporal inputs, and then apply those uncertainty estimates in computing modality-wise fusion weights. *In particular, we aim to estimate the aleatoric component of uncertainty [1] associated with different modalities for improved emotion recognition performance.* Unlike the epistemic component of uncertainty, which can be explained away with more data, aleatoric uncertainty captures noise or stochasticity that is inherent to an input signal. To give an example, in recognising emotional expressions from face images, the epistemic component can describe the uncertainty due to insufficient data for ‘happy’ class whereas the aleatoric component captures the uncertainty caused by factors like occluded facial regions, low-resolution face images, etc. In this work, we focus on estimating modality-wise aleatoric uncertainty in a multimodal emotion recognition model.

Being an intrinsically temporal and multimodal phenomenon, emotion recognition from multimodal inputs is a long-standing challenge in Affective Computing [2], [3], [4]. A meta-analysis

presented in [5] has shown that although emotion recognition can benefit from multimodal fusion in general, performance improvements are not significant when it comes to spontaneous emotions. We believe that uncertainty-aware multimodal fusion may have the potential to address this challenge, considering that the intensity of spontaneous emotions embedded in different modalities are likely to vary dynamically over time [6], [7].

Although Deep Neural Networks (DNNs) have been extensively applied to multimodal emotion recognition [8], [9], [10], [11], estimating modality-wise uncertainty for improved fusion performance is a relatively unexplored avenue. However, modelling predictive uncertainty (or confidence, its opposite) in DNNs received widespread attention in recent years [12], [13], [14], motivated by the observation that DNNs tend to make over-confident predictions [15], [16]. Most existing efforts towards uncertainty or confidence estimation in DNNs [13], [17] focus solely on reducing miscalibration errors, i.e., the mismatch between expected model estimation errors and their corresponding confidence scores. Recently, as an alternative perspective, Moon et al. [18] introduced the idea of learning to rank confidence scores for identifying the most reliable predictions.

In this work we argue that the estimated uncertainty scores must be simultaneously both *well-calibrated* and *well-ranked* (*ordinal*). The former is needed to accurately represent the correctness likelihood of a prediction for *an individual sample*. The latter is essential to effectively order predictions for *a group of samples* according to their correctness likelihoods. In other words, if an uncertainty estimate of an individual sample is well-calibrated, in the absence of its ground truth, the uncertainty score can serve as a proxy for its expected prediction error. If the uncertainty scores associated with different predictions are well-ranked or maintain ordinality, then one can use them to order their corresponding samples in terms of their reliability towards the target prediction, and to distinguish the most informative samples from the least informative samples.

For multimodal temporal learning, it is critical to estimate how informative the predictions made for different frames in different unimodal sequences are, towards estimating a common target label, so that the target-specific information can be reliably integrated [19]. In this work, we hypothesise that jointly learning these two properties – calibration and ordinality – can lead to more reliable uncertainty estimates for each modality, facilitating more effective uncertainty-weighted temporal context fusion. Based on this hypothesis, we propose an uncertainty modelling method that imposes the calibration and ordinality constraints jointly, as Fig. 1 illustrates.

For example, consider the task of classifying whether a person’s apparent emotional state as either ‘happy’ or ‘neutral’ by analysing a face image sequence and its speech signal. Assume that the face is covered with a mask in most frames, making the face modality less informative than the speech modality. In a unimodal setting, the face and speech classifiers are trained separately to output their corresponding ‘happy’ class probabilities. When well calibrated, these output probabilities should reflect the true correctness likelihoods of face and speech models’ predictions. Similarly, when constrained by ordinal ranking, the speech model’s output probability must be higher than the face

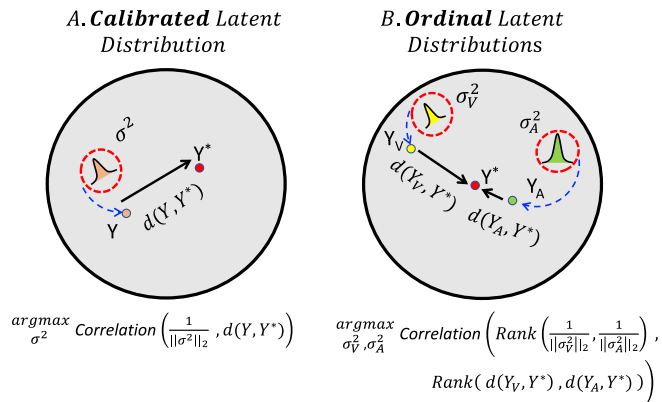


Fig. 1. Proposed latent distribution learning approach for multimodal fusion (Y_V and Y_A – unimodal predictions, Y^* – target label, and d – a distance function): The latent distributions’ variance values are learned to represent how informative the temporal context of each modality is, by applying the constraints: A. *Calibrated Latent Distribution*: Temporal context is modelled by a latent distribution learned under the *calibration constraint* so that its variance can act as a proxy for the target prediction error ($d(Y, Y^*)$). B. *Ordinal Latent Distributions*: The variance values of audio and visual temporal context distributions (σ_V^2 and σ_A^2) are learned under the *ordinal constraint* so that the audio and visual modalities are ranked according to their prediction errors ($d(Y_V, Y^*)$ and $d(Y_A, Y^*)$).

model’s probability, reflecting the relative uncertainty levels of the face and speech modalities w.r.t. each other.

In this work, we condition the unimodal latent distributions’ variance vectors such that they represent the information different modalities contain w.r.t. predicting emotion. The proposed method can be viewed as an uncertainty-aware extension of classical late fusion, but here the fusion is applied in the latent space of unimodal temporal context embeddings. This approach is different from a simple confidence-weighted late fusion model in which uncertainty is modelled directly over the unimodal output predictions.

In our proposed framework, denoted as Calibrated Ordinal Latent Distributions (COLD), we first learn the latent distributions (multivariate normal distributions) over the temporal context of audio and visual modalities separately, as Fig. 2 shows. We model the variance values of the audio and visual latent distributions, σ_V and σ_A , as the confidence measures towards emotion prediction. We design a novel training objective based on softmax distributional matching to encourage the variance norm values in each modality to be: (a) strongly correlated with the correctness likelihood of the unimodal predictions, and (b) ordinal in nature to effectively rank the relevance of different modalities towards emotion recognition. Thus, the calibrated and ordinal unimodal variance scores are learnt for effective uncertainty-weighted fusion, as shown in Fig. 2.

We evaluate the proposed COLD fusion approach on: (a) dimensional emotion recognition from face and voice modalities in the AVEC 2019 CES [4] and IEMOCAP [20] datasets, and (b) categorical emotion recognition from face, voice and text modalities in the CMU-MOSEI [21] and IEMOCAP datasets. Compared to the uncertainty-unaware fusion baselines, COLD fusion demonstrates noticeably better results on different multimodal emotion recognition tasks evaluated in this work.




Fig. 2. Overview of the proposed approach to an uncertainty-aware audiovisual fusion for emotion recognition: Modeling latent distributions over unimodal temporal context vectors to derive modality-wise uncertainty guided fusion weights. A detailed description of our approach is given in Section III.

For example, in dimensional emotion regression tasks COLD fusion shows $\sim 6\%$ average relative improvement over the best performing fusion baseline. Similarly, in the case of categorical emotion classification COLD fusion achieves $\sim 8.2\%$ relative accuracy improvement over the existing state-of-the-art model. Furthermore, we assess the robustness of different fusion models at test time by inducing noise into the visual modality through face masking. With the faces masked in 50% of the evaluation sequences, COLD fusion achieves $\sim 17\%$ average relative improvement over the best fusion baseline.

The key contributions of our work are as follows:

- We propose an uncertainty-aware multimodal fusion method that dynamically estimates the fusion weights to be assigned to unimodal features.
- We demonstrate how to jointly learn *well-calibrated* and *well-ranked* unimodal uncertainty estimates. For this purpose, we propose a simple softmax distributional matching loss function that applies to both regression and classification models.
- On both dimensional and categorical emotion recognition tasks, the proposed fusion approach shows noticeable performance gains and improved robustness to novel noise patterns encountered at test time.

II. RELATED WORK

Multimodal Affect Recognition: Humans rely primarily on visual (faces) and audio (voices) modalities to encode and express their affective or emotional states. Recognising dimensional emotions, valence (how pleasant an emotion is) and arousal (how active an emotion is), and categorical emotions (happy, sad, disgust, etc) from multiple modalities, is a widely studied problem in various prior works [7], [22], ranging from the almost a decade-long running annual AVEC challenge series [2], [3], [4] to the recently introduced MuSe challenge [23], [24], [25] and ABAW challenge [26], [27]. Beyond audiovisual modalities, some recent works (e.g. [28], [29]) explored how to fuse a wide range of contextual cues for reliably recognising expressed emotions.

We refer the reader to Poria et al. [30], Roust et al. [8], Jiang et al. [31] and Zhao et al. [32] for comprehensive surveys of affect recognition in multimodal settings and contemporary deep learning-specific advancements in it. Since our main focus in this work is on uncertainty-aware fusion models for emotion recognition, we review the literature closely related to the following key research topics: i) uncertainty modelling for emotion and expression recognition, ii) uncertainty-aware multimodal fusion, iii) calibrated uncertainty, and iv) ranking-based uncertainty.

Uncertainty Modelling for Emotion and Expression Recognition: In categorical facial expression recognition tasks, modelling predictive uncertainty is studied in several recent works [33], [34], [35], by estimating uncertainty in the space of low-dimensional feature embedding outputs from a Convolutional Neural Network (CNN) backbone. On the other hand, directly predicting emotion label uncertainty is explored in [36], but only in unimodal (video-only) settings. For uncertainty-aware multimodal emotion recognition, some prior works applied Kernel Entropy Component Analysis (KECA) [37] and Multi Modal-Hidden Markov Models (MM-HMMs) [38] by predicting modality-specific uncertainty measures for estimating the fusion weights.

Noting the limitations of deterministic function learning in DNNs for uncertainty modelling, Dang et al. [39] explored the application of Gaussian Process (GP) Regression to the fusion of emotion predictions. With the same motivation, in Affective Processes (APs) [40], [41], Neural Processes [42], [43] have been applied to the task of emotion recognition. By combining the abilities of GPs to learn function distributions with DNN's representation learning abilities, APs demonstrated superior generalisation performance over deterministic function learning models. Building on this idea of stochastic modelling of temporal functions, recently APs have been extended to multimodal settings in [44] based on a strictly model-based fusion approach, demonstrating impressive emotion recognition results. However, for uncertainty-aware temporal context modelling, APs heavily rely on the proxy labels predicted by a separate pre-trained backbone and a complex encoder-decoder formulation. In contrast to APs, our method aims to model the

temporal context uncertainty in a model-agnostic fashion, by just altering the output head of simple CNN+RNN models that are trained using some novel constrained optimisation objectives.

All the aforementioned methods demonstrated the potential of uncertainty-aware emotion recognition models over their uncertainty-unaware counterparts in general. However, they ignore two important aspects of uncertainty modelling: calibration and ordinality (ranking). In this work, we aim to demonstrate the significance of these two properties by hypothesising that learning well-calibrated and well-ranked uncertainty estimates is critical for improving multimodal emotion recognition performance.

Uncertainty-Aware Multimodal Fusion: For multimodal sensor fusion, several prior works [37], [45], [46], [47] explored uncertainty-aware or confidence-weighted averaging techniques for classic machine learning models before the advent of Deep Neural Networks (DNNs). Recently, Subedar et al. [48] applied Bayesian DNNs for uncertainty-aware audiovisual fusion to improve human activity recognition performance. Similarly, Tian et al. [49] explored the use of uncertainty estimation in fusing the softmax scores predicted using CNNs for semantic segmentation. Other notable approaches to uncertainty-aware multimodal fusion are based on optimal transport for cross-modal correspondence [50], random prior functions [51], boosted ensembles [52], and factorised deep markov models [53].

Although all the aforementioned methods demonstrated critical advantages over the models that predict only point estimates, they do not study the calibration properties of the estimated uncertainty scores. Further, such DNN models focus mainly on modelling absolute uncertainty estimates, whereas our focus is on jointly *learning the calibrated and relational uncertainty estimates* in an end-to-end fashion introducing a novel softmax distributional matching loss.

Calibrated Uncertainty: As DNNs tend to make overconfident predictions [15], [16], confidence calibration has received significant attention in recent years [15], [16]. Calibrating confidence or uncertainty estimates involves maximising the correlation between predictive accuracy values and predictive uncertainty scores. A wide variety of calibration techniques, particularly in classification settings, can be broadly categorised into explicit and implicit calibration categories [54]. In the former category, two types of post-hoc methods, binning-based and temperature-scaling, are applied to increase the reliability of DNN confidence estimates [13], [55]. In binning-based methods such as non-parametric histogram binning [56], calibrated confidence is estimated based on the average count of positive-class instances in each bin. This method is extended to jointly optimise the bin boundaries and their predictions in Isotonic Regression [57]. Temperature-scaling methods can be viewed as generalised versions of Platt scaling [58] using logistic regression for calibrating the class probabilities. We use temperature-scaling as a calibration baseline [13], [59] to compare against the uncertainty calibration performance of the proposed method, due to its simplicity.

Implicit calibration methods tailor the training objective of DNNs to minimise the prediction error and calibration error simultaneously. Addressing the limitations of standard

cross-entropy loss w.r.t. confidence calibration, various alternative loss functions such as focal loss [14], maximum mean calibration error [17], and accuracy vs uncertainty calibration [60], have been investigated recently. Calibrating regression models is relatively under-explored compared to the classification. Some recent works [61], [62], [63] made attempts to extend some of the aforementioned calibration techniques to continuous-valued predictions.

Ordinal or Ranking-based Uncertainty: In the existing uncertainty modelling works, the ordinal property of uncertainty estimates received less attention compared to the calibration property, which partly motivated the method introduced in this paper. Li et al. [64] proposed to model data uncertainty by inducing ordinality into probabilistic embeddings of face images. Towards uncertainty-aware regression problems, the results reported in [64] highlighted the key limitations of deterministic unordered embeddings compared to the probabilistic ordinal embeddings. Although not strictly ordinal, relative uncertainty modelling is explored for facial expression recognition in [34].

Other closely related works approached the problem of ordinal ranking of uncertainty estimates with different objectives such as failure prediction [65], out-of-distribution detection [66], and selective classification [67]. Fundamentally, all these objectives necessitate a method that can train the model to output well-ranked confidence or uncertainty scores. Among these existing methods, the one most closely related to ours is by Moon et al. [18], which proposes a Correctness Ranking Loss (CRL). CRL directly imposes ordinal ranking constraints on the confidence estimates of a DNN classifier. Similar to CRL, our proposed softmax distributional matching loss also constrains the ordinal-ranking property of uncertainty estimates. However, in addition to ordinal ranking, our method imposes the calibration property as well, most importantly by controlling the latent distribution variance, unlike in CRL. Moreover, our formulation generalises the idea of calibrated and ordinal uncertainty estimates to both classification and regression settings, using a common loss function computation.

III. MODEL-AGNOSTIC FUSION BASELINES

Before introducing our uncertainty-aware multimodal fusion formulation, we briefly discuss the general multimodal fusion techniques w.r.t. audiovisual emotion recognition and introduce the related notations. A fundamental question in multimodal learning concerns the optimal stage to perform fusion [68]. We consider the following three typical model-agnostic fusion methods as the baselines: feature fusion, temporal context fusion, and prediction fusion.

Preliminaries and Notations: As Fig. 2 illustrates, given a face video clip X_V with N frames and its corresponding speech signal X_A , using overlapping time windows, we first create N speech segments that correspond to the N visual frames. Here, we assume that both the signals X_V and X_A are annotated with a common dimensional emotion label, $Y^* = [Y_{valence}^*, Y_{arousal}^*]$ (either per-frame or per-sequence). We extract sequences of per-frame low dimensional features (Z_V, Z_A) from the face video and speech inputs using a two-stream network. This network

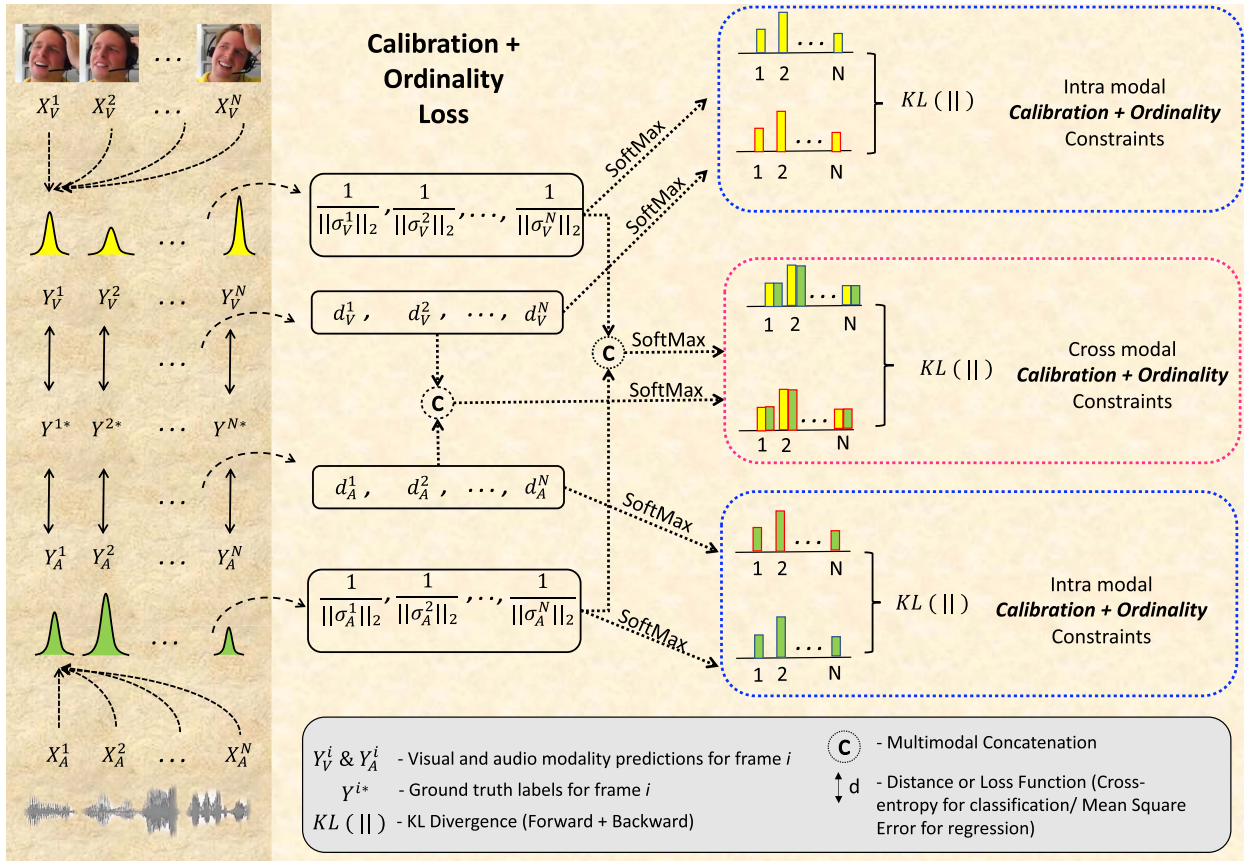


Fig. 3. COLD fusion training loss computation: To simultaneously impose the calibration and ordinality constraints on the unimodal latent distributions’ variance vectors, we minimise the softmax distributional matching loss (KL divergence) between the distance vectors $[d^i]$ and variance-norm vectors $[\frac{1}{\|\sigma^{i,2}\|_2}]$, in both intramodal and crossmodal settings.

is composed of a 2D CNN f_V and a 1D CNN f_A for processing the face images and speech segments respectively, $f_V : X_V \rightarrow [z_V^1, z_V^2, \dots, z_V^N]$ and $f_A : X_A \rightarrow [z_A^1, z_A^2, \dots, z_A^N]$. For unimodal emotion recognition, we process the temporal context from each modality separately from Z_V and Z_A using different temporal networks $g_V : Z_V \rightarrow Y_V$ and $g_A : Z_A \rightarrow Y_A$ to predict the emotion labels Y_V and Y_A .

Feature Fusion or early fusion integrates frame-level emotion cues present in the audiovisual features Z_V and Z_A (e.g., [69]), not accounting for commonly encountered temporal misalignment between different modalities [70]. Here, we concatenate the per-frame audiovisual features into a single sequence, $Z = [Z_V, Z_A]$, then pass it to a common temporal network $g_{AV} : Z \rightarrow Y$ to predict emotion labels.

Decision Fusion combines the unimodal emotion predictions Y_V and Y_A (e.g., [71]). Here, we apply predictive confidence based weighted averaging to perform the late fusion. Unlike early fusion, late fusion does not leverage the low-level correspondences among the emotion cues distributed over the audio and visual streams [68].

Temporal Context Fusion or simply context fusion integrates sequence-level emotion information aggregated in the form of audiovisual temporal context vectors h_V^i and h_A^i for frame i , produced by the temporal networks g_V and g_A respectively.

This method is also referred to as ‘feature fusion with RNNs’ or ‘mid-level’ fusion in some prior works [8], [72]. Note that here temporal context or simply context at i th frame refers to the emotion information present in frame i w.r.t. the emotion information carried by remaining frames in the input sequence. As a result, unlike early fusion, context fusion is bound to suffer less from the temporal misalignment between the emotion-related semantics of audio and visual feature sequences. Further, context fusion benefits from the low-level audiovisual correspondences in the emotion space, in contrast to late fusion.

Considering the above-mentioned critical advantages of temporal context fusion, in this work, we propose to learn an uncertainty-aware context fusion model for multimodal emotion recognition as discussed below.

IV. PROPOSED METHOD

Fig. 3 illustrates our proposed solution to uncertainty-aware multimodal fusion. Although this section describes the proposed fusion only in audiovisual settings, note that it can be easily extended to tasks with more than two modalities. In this section, we first discuss how we estimate modality-wise uncertainty by learning unimodal latent distributions over the temporal context, and we present our approach to derive the fusion weights based

on unimodal context variance. Then, we introduce two key optimisation constraints imposed on the variance norms of unimodal latent distributions and describe their implementations.

A. Uncertainty-Aware Audiovisual Context Fusion

Quantifying modality-wise uncertainty towards predicting a common target label is crucial to improve multimodal fusion performance. Our objective is to first quantify intramodal uncertainty in the temporal context space, and then use the estimated uncertainty scores to derive the fusion weights. To this end, we propose to learn unimodal latent distributions over the temporal context of the audio and visual modalities separately, as discussed below.

1) *Latent Distributions Over Unimodal Temporal Context:* Fig. 2 illustrates how we modify the temporal networks (Gated Recurrent Unit (GRU)-RNNs) g_V and g_A to output the parameters (mean and variance) of multivariate normal distributions $\mathcal{N}(\mu_V^i, \sigma_V^{i,2})$ and $\mathcal{N}(\mu_A^i, \sigma_A^{i,2})$ over the audio and visual temporal context vectors, respectively. Here, the term ‘temporal context’ refers to the hidden state outputs from the corresponding unimodal GRU blocks (g_A or g_V). For each modality separately, we learn this hidden state output as a multivariate normal distribution, instead of a typical deterministic embedding vector. We presume that these unimodal latent distributions are capable of representing modality-wise emotion information more effectively than deterministic embeddings.

Given a sequence of frames, $[X_1, X_2, \dots, X_T]$, in order to predict their corresponding target variables $[Y_1^*, Y_2^*, \dots, Y_T^*]$ it is important to learn the underlying temporal context information, which is a function of the frames present in the input sequence as well as the order in which they appear. By modelling the temporal context as a probability distribution, we propose to use the prediction error $\|Y_i - Y_i^*\|_2$ to constrain the contribution of each frame X_i in terms of its explained variance of the overall temporal context. Here, the idea of frame-wise explained variance of the temporal context refers to how much information a particular frame holds given all the rest of the frames, towards predicting the target variable Y_i^* . Thus, the higher the explained variance of a particular frame X_i , the more informative it is for accurately predicting the target variable.

Here our aim is to first estimate the informativeness of each modality towards the task of recognising emotions. To this end, we learn the temporal context variance such that it may represent how informative the temporal context of a particular modality is. For example, consider an audio-visual sequence in which all the audio frames have the same emotion (e.g., neutral tone), whereas the visual frames have more variations in terms of the emotional expressions. In this case, the fusion model must give more importance to the visual frames compared to the audio frames when predicting emotions. Guided by this intuition, our formulation aims to capture the emotion-related variance in the temporal context of each modality separately.

It is important to note the difference between the absolute variance of the temporal context distribution learned from all the frames and the explained temporal context variance of an individual frame. While the former can be thought of as a proxy

metric for uncertainty measurement, the latter can be viewed as a per-frame information metric w.r.t the target prediction. For the sake of simplicity, throughout this work, we use the term ‘context variance’ in order to refer to the explained variance of temporal context for a given frame in an input sequence. The above argument can be extended to a multimodal fusion setting as well, in which the explained temporal context variance of a particular modality can be used as a proxy for how informative that modality is w.r.t predicting a common target variable.

We model the variance of a unimodal latent distribution as a proxy for how informative that modality is w.r.t. predicting the target emotion, and we use the inverse of variance values to quantify how uncertain a particular modality is towards predicting emotion labels. Note that the potential of signal variance-based uncertainty modelling for multimodal fusion was already demonstrated in [73]. Similarly, learning latent distribution variance was determined to be capable of uncertainty modelling in [40]. Inspired by these ideas, we model the unimodal context variance norm values $\|\sigma_V^i\|_2$ and $\|\sigma_A^i\|_2$ to estimate how certain the audio and visual modalities are about predicting the emotion labels. Our approach to derive variance-based fusion weights for integrating the audiovisual information is discussed below.

2) *Context Distribution Variance-Based Fusion Weights:* For an input frame with index i , given its unimodal latent distributions $\mathcal{N}(\mu_V^i, \sigma_V^{i,2})$ and $\mathcal{N}(\mu_A^i, \sigma_A^{i,2})$ over its visual and audio temporal context embeddings separately, we first compute the L_2 norms of their variance values $\|\sigma_V^i\|_2$ and $\|\sigma_A^i\|_2$. As discussed above, these variance norm values are assumed to represent modality-specific certainty or informativeness w.r.t. predicting the target emotions. By normalising the variance norm values of the audio and visual modalities, we derive fusion weights that are used in a simple linear fusion model of the audiovisual temporal context (h_{VA}^i):

$$h_{VA}^i = w_V^i * h_V^i + w_A^i * h_A^i, \quad (1)$$

where h_V^i and h_A^i denote the visual and audio temporal context vectors, and w_V^i and w_A^i denote their corresponding weight values. The temporal context vectors h_V^i and h_A^i are sampled from their respective latent distributions, $h_V^i \sim \mathcal{N}(\mu_V^i, \sigma_V^{i,2})$ and $h_A^i \sim \mathcal{N}(\mu_A^i, \sigma_A^{i,2})$ during training. At test time, we set h_V^i and h_A^i to their corresponding mean vectors μ_V^i and μ_A^i for evaluation purposes.

Based on the unimodal context variance norm values ($\|\sigma_V^i\|_2$ and $\|\sigma_A^i\|_2$), the weight values w_V^i and w_A^i in (1) are computed as:

$$w_V^i = \frac{\|\sigma_V^i\|_2}{\left(\|\sigma_V^i\|_2 + \|\sigma_A^i\|_2\right)}, w_A^i = \frac{\|\sigma_A^i\|_2}{\left(\|\sigma_V^i\|_2 + \|\sigma_A^i\|_2\right)}. \quad (2)$$

Context variance modelling seems to be a simple yet effective approach to uncertainty-aware audiovisual fusion, yet learning audiovisual latent distributions with well-conditioned variance ranges is non-trivial in practice, as we show later in the experiments. To condition the variance values that can effectively capture intramodal uncertainty w.r.t. predicting the target labels,

we define a more principled model training that applies two key optimisation constraints: Calibration and Ordinality.

B. COLD: Calibrated and Ordinal Latent Distributions

To effectively learn the unimodal latent distributions for uncertainty-aware fusion, we propose to condition their variance values by applying optimisation constraints to the model training objective. We achieve this conditioning by imposing two key constraints: Calibration and Ordinality (or ranking) on the latent distribution variance vectors. When well-calibrated, an uncertainty score acts as a proxy for the correctness likelihood of its prediction for an individual input from a specific modality. In other words, well-calibrated uncertainty indicates the expected estimation error, i.e., how far the predicted emotion is expected to lie from its ground truth.

Given the predictions made for a set of frames from different modalities, when their uncertainty scores are well-ranked or maintain ordinality, we can effectively arrange the input unimodal frames according to their reliability for predicting a target emotion. In Fig. 1, we illustrate the definitions of both these constraints. *It is important to note the fundamental difference between these two constraints: while the calibration constraint is applied individually for each unimodal frame, the ordinality or ranking constraint is imposed jointly for a set of frames from different modalities.*

Calibration Constraint – this is imposed by regularising the unimodal context variance norms, $\|\sigma_V^i\|_2$ and $\|\sigma_A^i\|_2$, such that their values are strongly correlated with the correctness likelihood values of target emotion classes. In regression models, this constraint can be implemented by forcing the variance norm values to correlate with the Euclidean distance between their corresponding unimodal predictions Y_V and Y_A and their ground truth labels Y^* , as shown in Fig. 1. In other words, the context variance values are learnt as reliability measures indicating how far the emotion predictions are expected to lie from their ground truth labels. To impose this property on the variance values of both modalities, COLD fusion applies the following regularisation constraints,

$$\begin{aligned} \arg \max_{\sigma_V^2} \text{Correlation} \left(\frac{1}{\|\sigma_V^2\|_2}, d(Y_V, Y^*) \right) \\ \arg \max_{\sigma_A^2} \text{Correlation} \left(\frac{1}{\|\sigma_A^2\|_2}, d(Y_A, Y^*) \right) \end{aligned} \quad (3)$$

where $d(\cdot)$ denotes the distance function that measures the target emotion estimation error. Cross-entropy and Mean Squared Error (MSE) are used as the distance functions for the classification and regression models respectively.

Ordinality Constraint – this is applied to rank the frames of unimodal sequences, so that their uncertainty measures indicate how reliable different multimodal frames are w.r.t. each other. This ranking operation can be implemented as a simple ordering constraint which jointly regularises the unimodal context variance norm values, $\|\sigma_V^i\|_2$ and $\|\sigma_A^i\|_2$. Here, modality-wise reliability is again computed in terms of the distance values (see (3)) between different unimodal predictions and the ground

truth labels:

$$\begin{aligned} \arg \max_{\sigma_V^2, \sigma_A^2} \text{Correlation} \left(\text{Rank} \left(\frac{1}{\|\sigma_V^2\|_2}, \frac{1}{\|\sigma_A^2\|_2} \right), \right. \\ \left. \text{Rank} (d(Y_V, Y^*), d(Y_A, Y^*)) \right) \end{aligned} \quad (4)$$

1) Implementation: Calibration and Ordinality Constrained Training for Audiovisual Emotion Recognition: We train classification models of dimensional emotion recognition, in addition to the standard regression models used in the literature. In both cases, the underpinning principles of the COLD fusion are the same, but the training objective implementations differ slightly. To train the temporal context fusion models by imposing the calibration and ordinality constraints, we optimise the network to minimise a loss function composed of the following components:

Emotion Prediction Loss (L_{emo}) is computed using the standard cross-entropy function for training the classification models. For the regression models training, similar to [74], we use inverse Concordance Correlation Coefficient (CCC) loss (1.0 - CCC) in addition to MSE. This loss is computed for the predictions from unimodal (Y_V and Y_A) and multimodal (Y_{AV}) branches jointly (Fig. 2).

Calibration and Ordinality Loss (L_{CO}) combines the aforementioned constraints, defined in (3) and (4), into a single training objective using differentiable operations. Fig. 3 shows the steps involved in implementing this component: given an input sequence with N frames, we first compute their unimodal latent distributions followed by their corresponding unimodal predictions. To impose the calibration and ordinality constraints, we first compute two sets of vectors for each modality:

Distance Vectors: We collect the scalar distance values (d_V^i and d_A^i) between the unimodal predictions (Y_V^i and Y_A^i) and the ground truth labels (Y^{i*}) using either cross-entropy (classification) or MSE (regression) as the distance function. This step produces N-dimensional distance vectors, $D_V = [d_V^1, d_V^2, \dots, d_V^N]$ and $D_A = [d_A^1, d_A^2, \dots, d_A^N]$.

Variance-Norm Vectors: We collect the inverted unimodal context variance norm values into another set of N-dimensional vectors, S_V and S_A , as shown below:

$$\begin{aligned} S_V &= \left[\frac{1}{\|\sigma_V^1\|_2}, \frac{1}{\|\sigma_V^2\|_2}, \dots, \frac{1}{\|\sigma_V^N\|_2} \right] \\ S_A &= \left[\frac{1}{\|\sigma_A^1\|_2}, \frac{1}{\|\sigma_A^2\|_2}, \dots, \frac{1}{\|\sigma_A^N\|_2} \right]. \end{aligned} \quad (5)$$

Softmax Distributional Matching for Calibration and Ordinal Ranking: Note that the distance vectors and variance-norm vectors contain scalar values that summarise the properties of different embedding spaces, emotion labels, and temporal context, respectively. Hence, we assume that matching their properties by imposing the calibration and ordinality constraints directly in their original spaces, is not optimal. For this reason, as illustrated in Fig. 3, we first apply the softmax operation on the distance vectors and variance-norm vectors separately to generate the softmax distributions. Then, we impose the calibration and ordinality constraints by minimising the mismatch between softmax distributions of the variance-norm vectors and

distance vectors. This approach to calibration and ordinality loss computation based on soft-ranking is inspired by [75] in which softmax cross-entropy is used for ordinal regression.

As Fig. 3 shows, in both intramodal and crossmodal settings, we compute the softmax distributions of distance vectors (P_{D_V} , P_{D_A} , and $P_{D_{AV}}$) and variance-norm vectors (P_{S_V} , P_{S_A} , and $P_{S_{AV}}$). Note that in the crossmodal case, we first concatenate the audio and visual distance vectors and variance-norm vectors separately, i.e., $D_{AV} = [d_A^1, d_V^1, \dots, d_A^N, d_V^N]$ and $S_{AV} = [s_A^1, s_V^1, \dots, s_A^K, s_V^K]$. Then, we apply the softmax operation on the concatenated list which is $2N$ dimensional. Thus, the crossmodal softmax distributions capture the relative measures across both modalities. Now, to impose the calibration constraint, we minimise the KL divergence (both forward and backward) between the distance distributions and variance-norm distributions in both intramodal and crossmodal settings, as shown below:

$$L_{CO} = KL(P_D || P_S) + KL(P_S || P_D), \quad (6)$$

where P_D represents P_{D_V} and P_{D_A} , and P_S represents P_{S_V} and P_{S_A} in the intramodal loss computation. In the crossmodal case, P_D and P_S denote $P_{D_{AV}}$ and $P_{S_{AV}}$, respectively.

Variance Regularisation Loss (L_{regu}): Prior works [40], [76] on latent distribution learning in high-dimensional input spaces such as images, have reported that the variance collapse is a commonly encountered problem. Variance collapse occurs mainly because the network is encouraged to predict small variance σ^2 values to suppress the unstable gradients that arise while training the latent distribution models using Stochastic Gradient Descent. To prevent this problem, we include the regularisation term proposed in [76] in the training objective:

$$\begin{aligned} L_{regu} &= KL(\mathcal{N}(\mu, \sigma^2) || \mathcal{N}(\epsilon, \mathbf{I})) \\ &= -\frac{1}{2} (1 + \log \sigma^2 - \mu^2 - \sigma^2), \end{aligned} \quad (7)$$

where ϵ and \mathbf{I} denote the mean vector and an identity variance matrix respectively. Note that this regularisation term is applied to the audio and visual distributions, separately.

In summary, the COLD fusion training objective composed of the above-discussed loss components, is as follows:

$$\begin{aligned} L_{total} &= L_{emo} + \lambda_{CO_V} \cdot L_{CO_V} + \lambda_{CO_A} \cdot L_{CO_A} \\ &\quad + \lambda_{CO_{AV}} \cdot L_{CO_{AV}} + \lambda_R \cdot L_{regu}, \end{aligned} \quad (8)$$

where λ_{CO_V} (for visual-only), λ_{CO_A} (for audio-only), $\lambda_{CO_{AV}}$ (for audio and visual combined), and λ_R (for regularisation) are the optimisation hyperparameters that control the strength of each regularisation constraint.

V. EXPERIMENTS

We first discuss the details of dimensional and categorical emotion datasets used for evaluating the proposed COLD fusion model. Detailed information about each dataset can be found in [4], [20], [21]. Then, we discuss the regression and classification formulations of emotion recognition and the evaluation metrics used for dimensional and categorical emotion tasks, along with a standard uncertainty calibration error metric that applies to the classification models. Finally, we present the details of the

network architectures, fusion model implementations, and their optimisation.

A. Datasets

1) *Dimensional Emotion Recognition:* For spontaneous dimensional emotion recognition, we used the AVEC 2019 CES challenge corpus [4] which is designed for in-the-wild emotion recognition in cross-cultural settings as part of the SEWA project [77]. This corpus is composed of 8.5 hours of audiovisual recordings collected from German, Hungarian, and Chinese participants. All videos in this corpus are annotated with continuous-valued valence and arousal labels in the range $[-1, 1]$. Note that the train and validation partitions are composed of only German and Hungarian cultures. As the labels for the test set (which has the Chinese culture in addition) are not publicly available, we report results on the validation set.

For *Acted* emotion recognition, the Interactive Emotional Dyadic Motion Capture (IEMOCAP) dataset [20] is used. This dataset constitutes 12 hours of audiovisual data annotated with utterance-level labels of valence and arousal. Here, we normalised the original emotion labels to the range $[-1, 1]$. Among the available five sessions in this corpus, we used the first four sessions' data for training. Note that the COLD fusion model training involves tuning of multiple regularisation constraints ((8)). Thus, the usual 5-fold cross-validation evaluation is found to be computationally expensive as it requires the values of λ_{CO_V} , λ_{CO_A} , $\lambda_{CO_{AV}}$, and λ_R to be tuned for every fold. For this reason, we used the speaker-independent partitions of the fifth session as validation and test sets, the same as the first fold's validation and test sets used in the existing works (e.g. [78], [79]) that apply 5-fold cross-validation.

On both the emotion datasets we trained and evaluated our audiovisual fusion models in regression as well as classification settings. To train the regression models, we directly used the continuous-valued labels as targets in the range $[-1, 1]$. For classification, we first mapped the continuous emotion values to three different classes for valence (positive, neutral, negative) and arousal (high, neutral, low) individually. For this binning, we chose the thresholds of -0.05 and 0.05 to draw the boundaries between the three above-mentioned bins. We adjusted the binning thresholds and picked the aforementioned values, to minimise the imbalances in the resultant class-wise label distributions.

Addressing Imbalanced Emotion Class Label Distributions: Despite carefully tuning the binning thresholds, class-wise label distributions of the dimensional emotion datasets still have significant imbalances, as shown in Fig. 4. To mitigate the effect of this problem, we applied two general techniques while training the classification models: a. non-uniform sampling of the training instances for different classes and b. class-weighted cross-entropy loss. In the former, we modified the sampling criteria to over sample for the minority classes and under sample for the majority classes based on the number of examples available for each class in the train set. In the latter technique, we divided the cross-entropy loss values for different classes by their relative bin size (in the train set).

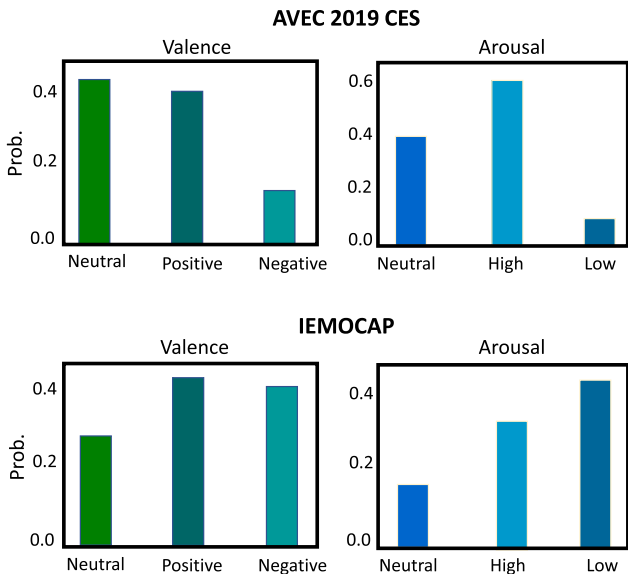


Fig. 4. Class imbalances in the distributions of valence and arousal labels prepared for 3-way classification on the AVEC 2019 CES and IEMOCAP datasets.

2) *Categorical Emotion Recognition*: For spontaneous categorical emotion recognition, we used the CMU-MOSEI dataset [21], a large-scale dataset for in-the-wild multimodal emotion recognition. This dataset is composed of 23,453 video utterances collected from YouTube monologues of 1000 distinct speakers. Each utterance contains three modalities: image sequences sampled at 30 Hz, audio waveforms with a sample rate of 44.1 kHz, and their corresponding text transcripts. All the utterances are manually annotated with 6 categorical emotions: *angry*, *disgust*, *fear*, *happy*, *sad*, and *surprise*. Here, we used the same training, validation and test partitions that are provided as part of the CMU-Multimodal Software Development Kit.¹

For acted categorical emotion recognition, we used the IEMOCAP dataset with the labels of six basic emotions: *neutral*, *angry*, *happy*, *sad*, *excited* and *frustrated*. Following the existing works [80], [81], we used a pre-processed version of this dataset that contains 7,380 utterances, in which each utterance contains an image sequence sampled at 30 Hz, an audio waveform sampled at 16 kHz, and its text transcript. We followed the same training (70%), validation (10%) and test (20%) splits used in prior works (e.g. [80]).

B. Evaluation Metrics

Regression models' performance is measured using Lin's Concordance Correlation Coefficient (CCC) [82] between the predicted emotions y^o and their ground truth labels y^*

$$CCC = \frac{\rho_{y^*y^o} \cdot \sigma_{y^*} \cdot \sigma_{y^o}}{(\mu_{y^*} - \mu_{y^o})^2 + \sigma_{y^*}^2 + \sigma_{y^o}^2}, \quad (9)$$

where $\rho_{y^*y^o}$ denotes the Pearson's coefficient of correlation between y^* and y^o , and (μ_{y^*}, μ_{y^o}) and $(\sigma_{y^*}, \sigma_{y^o})$ denote their mean and standard deviation values, respectively.

Classification models of dimensional emotions are evaluated using precision, recall, and F1 score. Given the imbalanced emotion class distributions (see Fig. 4), for these three metrics we report unweighted or macro averaged values of the three emotion classes, so that the average values are not biased towards the most dominant classes. For evaluating the categorical emotion models, following prior works [80], [81], [83], [84], we used (a) the accuracy and F1 score metrics for IEMOCAP and (b) the weighted accuracy and F1 score for CMU-MOSEI.

Uncertainty Calibration Errors of the classification models are measured to analyse the deviations between the true class likelihoods p and the predicted class confidence estimates \hat{p} . Reliability diagrams [13] are used as empirical approximations to visually represent the confidence calibration errors. For plotting these diagrams, first, the accuracy and confidence axes are binned into equally-sized intervals and then, for each interval mean accuracy values are plotted against their corresponding mean confidence scores. For a perfectly calibrated model, the reliability diagram is supposed to be an identity function, i.e., accuracy and confidence should have the same values. Expected Calibration Error (ECE), a scalar summary statistic of the reliability diagram, computes the weighted average of calibration errors over all the intervals in a reliability diagram.

$$ECE = \sum_{m=1}^M \frac{|I_m|}{N} |Acc(I_m) - Conf(I_m)|, \quad (10)$$

where I_m denotes the m th interval, M is the total number of intervals, and N is the total number of samples.

C. Network Architectures

1) *Feature Extraction for Dimensional Emotion Models*: *Visual CNN Backbone*: EmoFAN [85], a 2D CNN proposed recently for facial feature extraction, is proven highly efficient by building on hour-glass-based network architectures. This CNN backbone, pretrained on 2D face alignment task, has been found very efficient for transfer learning tasks [86], [87]. We used its pretrained model² on image-based emotion recognition on the AffectNet dataset [88]. Using this backbone, we extracted a 512D feature vector per frame.

Audio CNN Backbone: We adopted a 2D CNN backbone proposed in [89] for extracting speech signal features in an end-to-end fashion. Here, we applied a VGGish [90] pre-trained module to 2D Mel-spectrograms that are derived by setting the hop size and window length values to 0.1 s and 1 s respectively. Similar to [89], we fine-tuned only the last two fully connected layers of this VGGish module. To differentiate the interlocutor's information from that of the target speaker, we implemented the feature dimensionality-doubling technique proposed in [91].

Data Augmentation: We applied strong data augmentation techniques to the audiovisual inputs to minimise the overfitting problem. It is important to note that under heavy overfitting, the COLD loss function ((6)) may collapse since the calibration and ordinality constraints rely on the prediction errors of the training

¹<https://github.com/A2Zadeh/CMU-MultimodalSDK>

²Pretrained models of Toisoul et al. [86] are available at <https://github.com/face-analysis/emonet>

instances. For face image data, we applied horizontal flipping with the probability set to 0.5, random scaling by a factor of 0.25, random translation by +/- 30 pixels, and random rotation by 30°. In the audio case, we applied SpecAugment [92] which directly augments the 2D spectrogram itself, instead of its original 1D waveform. Here, we applied the standard SpecAugment operations: time warping, frequency masking and time masking, with their order defined arbitrarily. The parameters³ of time warping (ω), frequency masking (f), and time masking (t) are chosen from different uniform distributions in the range [0, 50], [0,27], and [0,40] respectively.

2) *Feature Extraction for Categorical Emotion Models*: Following the existing works [80], [81], we applied the early-stage feature extraction on the aligned multimodal data. The visual features containing 35 facial action units are extracted using Facet.⁴ The audio features extracted using COVAREP [93] contain glottal source parameters, Mel-frequency cepstral coefficients, etc. Similar to the prior works [80], [81], we used 74-dimensional and 144-dimensional audio features for CMU-MOSEI and IEMOCAP datasets, respectively. The text feature vectors with 300 dimensions are prepared by tokenising the text data at word level and then extracting their GloVe [94]⁵ embeddings.

3) *Temporal Networks*: In dimensional emotion recognition models, the temporal networks are stacked on top of the unimodal CNN backbones to model the temporal dynamics and integrate the multimodal affect information. Note that all the fusion models evaluated in this work follow different temporal network implementations. However, all the temporal networks have the following GRU block in common: a 2-layer bidirectional GRU module followed by a fully connected (FC) output layer. This GRU block contains 256 hidden units with the dropout value set to 0.5. The number of GRU blocks and their input-output dimensionality vary across different fusion models, as discussed below.

In feature fusion, a single GRU+FC block is used to process the input feature sequence that is prepared via frame-wise concatenation of the unimodal embeddings, whereas, in the prediction fusion, different unimodal temporal models (GRU+FC) are applied separately, and their output softmax label distributions are aggregated into the final predictions. The context fusion implementation has two different GRU blocks, but a common FC layer. As shown in Fig. 2, COLD fusion is similar to the context fusion, but with the GRU block’s output layer modified to predict the mean and variance vectors. Note that we trained the unimodal output branches simultaneously along with the fusion branch in all the multimodal models (see Fig. 2).

In categorical emotion recognition models, the pre-extracted visual, audio, and text features are directly fed into their corresponding temporal networks, which are composed of the same GRU+FC blocks used in the dimensional emotion models. Except for the number of input units, which depend on the

³ ω – warping length, f – number of consecutive Mel frequency channels masked, t – number of consecutive time steps masked

⁴iMotions. Facial expression analysis, 2017.

⁵glove.840B.300 d: <https://nlp.stanford.edu/projects/glove/>

TABLE I
DIMENSIONAL EMOTION REGRESSION RESULTS ON THE AVEC 2019 CES VALIDATION SET (CCC: CONCORDANCE CORRELATION COEFFICIENT)

Model	Valence CCC \uparrow	Arousal CCC \uparrow	Avg. CCC \uparrow
AVEC CES Winners:Aud [98]	0.388	0.518	0.453
Aud-branch	0.369	0.465	0.417
AVEC CES Winners:Vis [98]	0.579	0.594	0.586
Vis-branch	0.511	0.514	0.512
AVEC CES Winners:AV Fusion [98]	0.614	0.645	0.629
AV Feature Fusion	0.515	0.509	0.512
AV Prediction Fusion	0.552	0.617	0.584
AV Context Fusion	0.578	0.620	0.599
AV COLD Fusion	0.611	0.661	0.636

input feature dimensionality, all the network parameters are the same in both the cases. In the COLD fusion module, due to the presence of the third modality (i.e., text features) present in the categorical emotion models, (2) is modified to accommodate three modalities and the calibration and ordinal constraints, (3) and (4), are modified to compute pair-wise correlations for the six possible combinations of the audio, visual and text modalities.

D. Optimisation Details

The batch size, learning rate, and weight decay values chosen for training all these models are 4, 5e-3, and 1e-4, respectively. For tuning the learning rate, we used Cosine annealing coupled with warm restarts [95] (the number of epochs for the first restart set to 1 and the multiplication factor set to 2). We used Adam optimiser [96] for training all the models evaluated in this work.

For dimensional emotion recognition, we used input sequences of 30 seconds duration with per-frame and per-sequence targets on the AVEC 2019 and IEMOCAP datasets respectively. The visual and audio backbones and all the fusion models are trained by jointly minimising the CCC loss [74] and mean squared error for the regression task and class-weighted cross-entropy loss for the classification task. For finding the optimal values of hyper-parameters, we used the IEMOCAP validation set and the same optimal values are applied to the models trained on the AVEC 2019 corpus. The hyper-parameter values in the loss function ((6)) are tuned on the logarithmic scale in the range [1e-5, 1e+5] using RayTune [97]. Based on the IEMOCAP validation set performance, the following values are found to be optimal: 1e-3 for λ_{CO_V} , λ_{CO_A} and $\lambda_{CO_{AV}}$, and 1e-4 for λ_R . We applied the same hyperparameter values to the models trained on the AVEC 2019 corpus as well.

For categorical emotion recognition, we used sequences of 100 frames. The temporal networks are trained using the standard cross-entropy loss. The hyper-parameters are tuned separately on the validation sets of CMU-MOSEI and IEMOCAP. The following values are found to be optimal: 1e-2 and 5e-3 for $\{\lambda_{CO_V}, \lambda_{CO_A}, \text{ and } \lambda_{CO_{AV}}\}$ on CMU-MOSEI and IEMOCAP respectively and, 1e-4 and 5e-5 for λ_R on CMU-MOSEI and IEMOCAP respectively.

TABLE II
DIMENSIONAL EMOTION REGRESSION RESULTS ON THE IEMOCAP TEST SET
(CCC: CONCORDANCE CORRELATION COEFFICIENT)

Model	Valence	Arousal	Avg.
	CCC \uparrow	CCC \uparrow	CCC \uparrow
Aud-branch	0.694	0.453	0.573
Vis-branch	0.496	0.355	0.425
AV Feature Fusion	0.568	0.390	0.479
AV Prediction Fusion	0.696	0.481	0.578
AV Context Fusion	0.690	0.487	0.588
AV COLD Fusion	0.723	0.504	0.613

VI. RESULTS AND DISCUSSION

We first present the results of dimensional and categorical emotion recognition models based on different audiovisual fusion techniques. By inducing visual noise through face masking, we investigate the robustness of the proposed COLD fusion compared to the standard fusion baselines. Then, we analyse the uncertainty calibration performance of the COLD fusion model, particularly in classification settings. Finally, a qualitative analysis of modality-wise fusion weights is presented to demonstrate the calibration and ordinal ranking proprieties of the COLD fusion model.

A. Dimensional Emotion Recognition Results

Regression performance of different unimodal (Aud-branch and Vis-branch) and multimodal (AV) predictions are presented in Tables I and II for the AVEC 2019 CES (spontaneous emotion recognition) and IEMOCAP (acted emotion recognition) corpora, respectively. In both cases, COLD fusion consistently outperformed the standard fusion baselines (feature, prediction and context) as well as the unimodal results. When compared to the best performing CNN+RNN fusion baselines COLD fusion achieved $\sim 6\%$ average relative improvement.

Compared to the winners of the AVEC 2019 challenge, Zhao et al. [98], COLD fusion performs well in terms of arousal and mean CCC scores. However, it is slightly worse in the case of valence CCC. Note that Zhao et al. [98] use a domain adaptation technique to cope with the cross-cultural variations in audiovisual emotion expressions. However, our focus is not on coping with the cross-cultural variations, but primarily on improving the fusion performance. It is important to note that our fusion technique is, in principle, complementary to the domain adaptation used in [98]. More advanced temporal models such as Affective Processes [40], [41], [44] demonstrated superior generalisation performance than the RNNs in recent years. However, since this work mainly focuses on capturing temporal uncertainty for model-agnostic fusion based on simple CNN+RNN formulations, such complex temporal models based on APs are not included in this comparison, to not clutter the analysis of standard model-agnostic fusion methods presented here.

In Appendix A, available online, we compare the proposed COLD fusion and a multimodal Transformer baseline [99] on

TABLE III
DIMENSIONAL EMOTION 3-WAY CLASSIFICATION RESULTS (P: PRECISION, R: RECALL, F1: F1 SCORE) ON THE AVEC 2019 CES VALIDATION SET

Model	Valence			Arousal		
	P \uparrow	R \uparrow	F1 \uparrow	P \uparrow	R \uparrow	F1 \uparrow
Aud-branch	68.3	48.2	56.6	74.3	50.2	59.9
Vis-branch	70.9	58.1	63.9	76.8	70.3	73.4
AV Feature Fusion	67.8	60.2	63.8	73.4	68.2	70.7
AV Prediction Fusion	68.9	60.5	64.4	77.0	69.4	73.0
AV Context Fusion	75.0	60.6	67.0	77.1	71.1	73.9
AV COLD Fusion	76.8	62.4	68.9	79.5	74.0	76.5

TABLE IV
DIMENSIONAL EMOTION 3-WAY CLASSIFICATION RESULTS (P: PRECISION, R: RECALL, F1: F1 SCORE) ON THE IEMOCAP TEST SET

Model	Valence			Arousal		
	P \uparrow	R \uparrow	F1 \uparrow	P \uparrow	R \uparrow	F1 \uparrow
Aud-branch	55.5	55.3	55.4	62.4	55.1	58.5
Vis-branch	53.7	48.7	51.0	56.8	49.2	52.7
AV Feature Fusion	62.8	47.8	54.3	58.4	54.2	56.2
AV Prediction Fusion	64.5	57.3	60.7	58.9	57.0	57.9
AV Context Fusion	64.3	61.6	62.9	61.8	55.6	58.6
AV COLD Fusion	66.3	61.5	63.8	64.9	59.5	62.1

the AVEC 2019 dimensional emotion regression task. Here also, COLD fusion clearly outperformed the transformer baseline by a noticeable margin, especially in arousal prediction.

Appendix B, available online presents an ablation study of different components in the COLD fusion formulation, by nullifying different hyperparameters to modify the COLD training objective ((8)). These results, as shown in Table 11, show the importance of calibration, ordinal, and variance regularisation constraints to the overall performance improvements achieved by COLD fusion. In Appendix C, available online, we present the results of statistical significance tests, further validating the improvements achieved by COLD fusion over the standard fusion baselines.

Classification performance on the AVEC 2019 CES and IEMOCAP corpora is presented in Tables III and IV. Similar to the regression results, COLD fusion demonstrates superior emotion classification results on both datasets. Note that here, we pose the original regression problem as a 3-way classification problem by discretising the continuous emotion labels. For this reason, we do not have any existing benchmarks for comparison in this particular classification setting. Nevertheless, the performance improvements achieved by the COLD fusion are consistent for both valence and arousal in terms of all three metrics, except for the valence recall on IEMOCAP.

Unimodal Performance Analysis: It is interesting to note that in the AVEC 2019 case, the visual modality (Vis-branch) has a considerably better performance compared to the audio modality (Aud-branch), while it is vice versa in the case of the IEMOCAP dataset. This discrepancy may be due to the difference in the quality of the video data in terms of face image

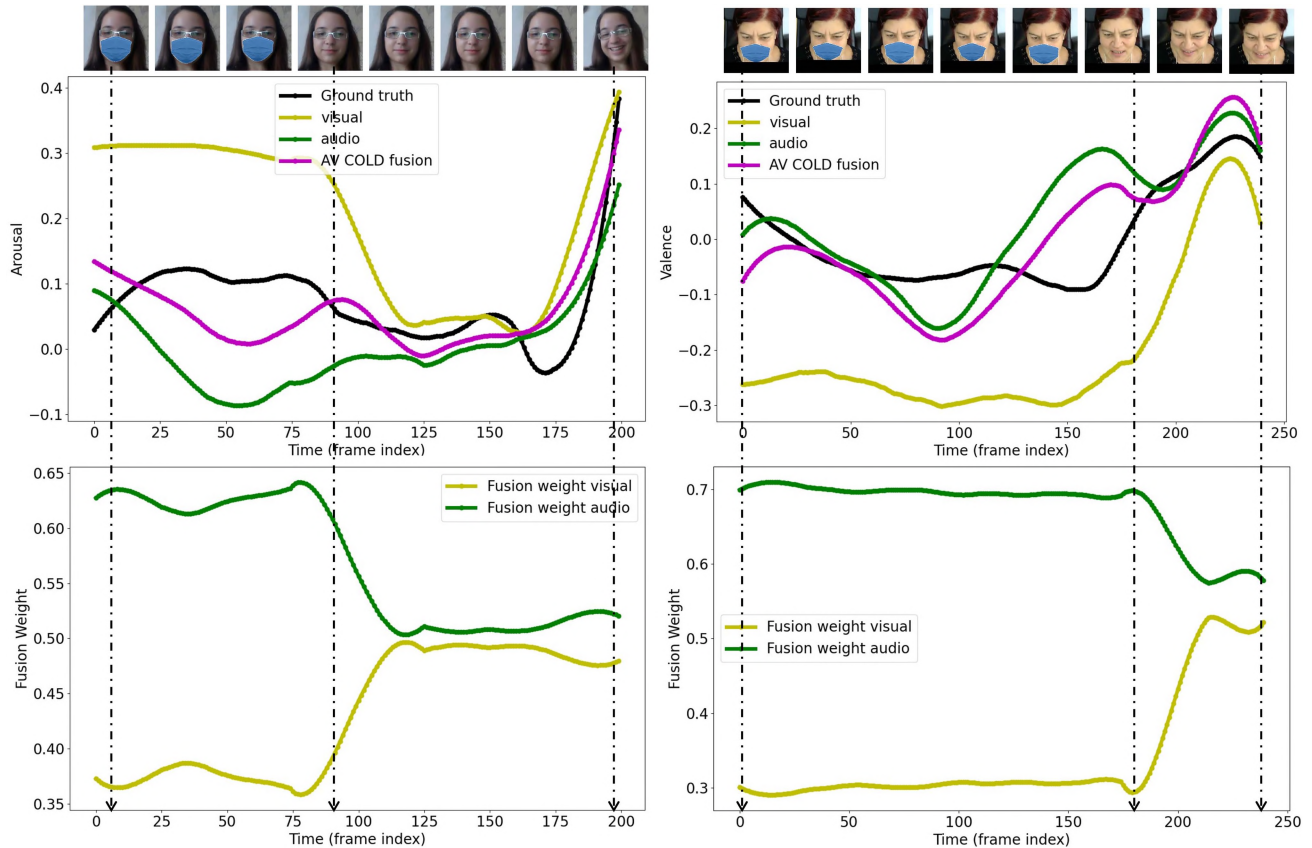


Fig. 5. Dynamic adaptation of COLD fusion weights when presented with novel noise patterns induced into the visual inputs: At test time, face masking is applied to randomly chosen consecutive frames in the AVEC 2019 CES validation examples. When the visual modality is noisy, i.e., containing faces with masks, AV COLD fusion output relies more on the audio modality (note the gaps between visual predictions and AV COLD fusion predictions, and modality-wise fusion weights). After removing the face masks, the fusion weight values adapt accordingly, hence, the fusion outputs.

resolution. Despite such dataset-specific differences, our COLD fusion technique shows consistent performance improvements in the multimodal classification and regression settings for both datasets.

Analysis of Fusion Baselines: Among the fusion methods that we evaluated here, temporal context or simply context fusion is found to be the second-best performing method after the proposed COLD fusion, on both datasets. Note that here, the temporal context refers to the output of the unimodal GRU block, and unimodal predictions are generated by applying a shallow fully connected network to the unimodal context vector. Thus, the context vectors can be viewed as higher-dimensional descriptors of the final unimodal predictions. Based on this assumption, in theory, the performance of context fusion is bound to be either better or at least as good as the prediction fusion, justifying the trends observed in our experimental results.

We notice that the feature fusion performance is inferior to all the remaining fusion techniques, and prediction fusion performs better than feature fusion. This result is consistent with an observation that prediction fusion achieves better results compared to feature fusion in general, as reported in the existing multimodal affect recognition literature [71]. It is worth noting that the results of feature fusion are worse than that of the best performing unimodal models on both datasets, i.e., the visual

(Vis-branch) model on AVEC 2019 and the audio (Aud-branch) model on IEMOCAP. This performance degradation may be due to not explicitly correcting the temporal misalignment effects [70], which are heuristically derived in general [4]. This result indicates that integrating multimodal emotion information at the feature-level or frame-level could be suboptimal most likely due to the temporal misalignment issues, given that continuous emotion information is expressed in the audiovisual modalities at different frame rates [8], [72].

Dynamic Adaptation of Fusion Weights in the Presence of Noise: In this experiment, we aim to understand how different fusion models perform when presented with novel noise patterns at test time. By inducing noise into the visual modality through face masking, here, we investigate the performance of different fusion baselines in comparison with the COLD fusion. For this evaluation, we overlaid the face masks as external occlusions on the image sequences using the method proposed in MaskTheFace [100].⁶ We applied MaskTheFace to 50% of the randomly chosen consecutive frames of the AVEC 2019 CES validation set sequences, as shown in Fig. 5. Note that all the fusion models evaluated here have not seen faces with masks during their training. As Table VII shows, in this

⁶<https://github.com/aqeelanwar/MaskTheFace>

TABLE V

DIMENSIONAL EMOTION CLASSIFICATION *CALIBRATION* RESULTS ON THE **AVEC 2019 CES VALIDATION SET** (ECE: EXPECTED CALIBRATION ERROR, BTS: BEFORE TEMPERATURE SCALING, ATS: AFTER TEMPERATURE SCALING)

Model	Valence ECE ↓		Arousal ECE ↓	
	BTS	ATS	BTS	ATS
Aud-branch	13.6e-2	6.3e-2	3.2e-2	2.8e-2
Vis-branch	8.9e-2	7.1e-2	12.6e-2	3.1e-2
AV Feature Fusion	6.1e-2	5.0e-2	5.5e-2	3.3e-2
AV Prediction Fusion	8.7e-2	5.1e-2	2.5e-2	2.6e-2
AV Context Fusion	6.9e-2	4.0e-2	6.3e-2	3.0e-2
AV COLD Fusion	3.7e-2	4.3e-2	1.3e-2	0.9e-2

TABLE VI

DIMENSIONAL EMOTION CLASSIFICATION *CALIBRATION* RESULTS ON THE **IEMOCAP TEST SET** (ECE: EXPECTED CALIBRATION ERROR, BTS: BEFORE TEMPERATURE SCALING, ATS: AFTER TEMPERATURE SCALING)

Model	Valence ECE ↓		Arousal ECE ↓	
	BTS	ATS	BTS	ATS
Aud-branch	7.4e-2	6.0e-2	13.2e-2	8.1e-2
Vis-branch	8.0e-2	4.1e-2	7.4e-2	5.0e-2
AV Feature Fusion	7.5e-2	5.6e-2	7.0e-2	5.8e-2
AV Prediction Fusion	8.4e-2	6.4e-2	8.4e-2	4.2e-2
AV Context Fusion	9.7e-2	6.6e-2	7.9e-2	6.0e-2
AV COLD Fusion	6.1e-2	3.8e-2	1.1e-2	2.0e-2

TABLE VII

IMPACT OF VISUAL NOISE (EXTERNAL OCCLUSIONS) ON THE AV FUSION MODELS: DIMENSIONAL EMOTION *REGRESSION* RESULTS WITH 50% OF RANDOMLY CHOSEN FACE IMAGES MASKED DURING EVALUATION (SEE FIG. 5) ON THE AVEC 2019 CES VALIDATION SET

Model	Valence	Arousal	Avg.
	CCC ↑	CCC ↑	CCC ↑
AV Feature Fusion	0.378	0.351	0.364
AV Prediction Fusion	0.363	0.545	0.454
AV Context Fusion	0.385	0.508	0.445
AV COLD Fusion	0.491	0.574	0.528

noise-induced evaluation setup, performance drop compared to the noise-free evaluation (Table I) is considerably higher for all three fusion baselines (feature, prediction, and context) than for the COLD fusion. Furthermore, the relative performance difference between the COLD fusion and the best-performing fusion baselines is increased from $\sim 6\%$ in noise-free settings to $\sim 17\%$ in this noise-induced case.

Fig. 5 compares the COLD fusion predictions with the predictions from visual and audio branches, along with the inferred modality-wise fusion weight scores. We can clearly see that the visual fusion weights are much lower for the frames with masks compared to the frames without masks, and as a result, the final predictions rely more on the audio modality in the presence of visual noise. This result demonstrates the ability of COLD fusion to dynamically adjust the importance of a specific modality according to its informativeness towards recognising the target emotions.

B. Categorical Emotion Recognition Results

The comparative results for the categorical emotion recognition tasks are presented in Table VIII (CMU-MOSEI) and Table IX (IEMOCAP). This comparison considers the following baselines: late fusion models based on LSTMs and Transformers, existing multimodal benchmarks and a SOTA model (AMOA [83]) among the two-phase models. Note that most existing models evaluated on CMU-MOSEI and IEMOCAP take a two-phase approach to multimodal emotion recognition, in which unimodal hand-crafted feature extraction and multimodal temporal fusion are performed separately. In line with those works, we evaluated the COLD fusion using the same two-phase architecture.

As shown in Tables VIII and IX, COLD fusion achieves new SOTA performance among the two-phase models. Note that on the both datasets there is a noticeable performance difference between the context fusion and COLD fusion models, which demonstrates the importance of the proposed calibration and ordinal constraints on temporal latent distribution learning. On CMU-MOSEI, compared to the existing SOTA (AMOA), COLD fusion achieves 8.2% and 1.8% relative improvements in terms of the average weighted accuracy and F1 scores respectively. On IEMOCAP, COLD fusion demonstrates the best accuracy and the second best F1 score. Here, the model with the highest F1 score is based on a multimodal transformer (Mult [99]), whereas the COLD fusion model implemented in this work uses GRUs for modelling the temporal dynamics. For further performance improvements, the proposed COLD fusion model can be integrated with transformer-based temporal models for combining the best of both worlds.

Compared to the two-phase models considered here for evaluation, some recently proposed fully end-to-end models such as the ones in [80], [84], [101] demonstrated improved emotion recognition performance but at the cost of significantly increased model training complexity. Although the COLD fusion framework is not evaluated in such models in this work, its ability to achieve robust multimodal fusion can be extended to fully end-to-end models as well for additional performance gains.

To demonstrate the applicability of COLD fusion to other multimodal tasks, besides emotion recognition, we evaluated it on utterance-level multimodal (AVL) sentiment analysis tasks on the CMU-MOSEI dataset. Refer to Appendix E, available online for the sentiment classification and regression results of the COLD fusion model in comparison with the existing baselines. In this case, COLD fusion achieves competitive results compared to the best-performing baseline (MISA [102]) and it shows the best results when coupled with MISA.

Overall, the multimodal categorical emotion and sentiment recognition results demonstrate the importance of learning well-calibrated and well-ranked uncertainty scores for improved multimodal fusion performance. These experiments also show that the COLD fusion formulation can be easily extended to models with more than two modalities.

C. Uncertainty Calibration Performance Analysis

To measure the quality of uncertainty estimates, we computed Expected Calibration Error (ECE) (see Section V-B) values for

TABLE VIII
CATEGORICAL EMOTION CLASSIFICATION RESULTS ON THE **CMU-MOSEI** TEST SET (WAC. AND F1 INDICATE WEIGHTED ACCURACY AND F1 SCORES RESPECTIVELY AND [†] INDICATES THE BASELINE MODELS' RESULTS FROM LI ET AL. [84])

Model	Angry		Disgusted		Fear		Happy		Sad		Surprised		Average	
	WAc.↑	F1↑	WAc.↑	F1↑	WAc.↑	F1↑	WAc.↑	F1↑	WAc.↑	F1↑	WAc.↑	F1↑	WAc.↑	F1↑
LSTM Pred. Fusion [†]	64.5	47.1	70.5	49.8	61.7	22.2	61.3	73.2	63.4	47.2	57.1	20.6	63.1	43.3
Transf. Pred. Fusion [†]	65.3	47.7	74.4	51.9	62.1	24.0	60.6	72.9	60.1	45.5	62.1	24.2	64.1	44.4
EmoEmbs [81]	66.8	49.4	69.6	48.7	63.8	23.4	61.2	71.9	60.5	47.5	63.3	24.0	64.2	44.2
MuT [99] [†]	64.9	47.5	71.6	49.3	62.9	25.3	67.2	75.4	64.0	48.3	61.4	25.6	65.4	45.2
AMOA [83] [†]	66.4	47.5	74.9	52.2	62.0	25.1	62.6	73.4	63.8	47.2	64.3	26.5	65.7	45.3
AVL Context Fusion	61.8	47.3	79.6	50.6	64.9	22.7	66.3	73.3	55.8	48.1	63.4	24.5	65.3	44.5
AVL COLD Fusion	69.7	49.3	77.6	52.9	79.8	25.6	65.8	74.2	54.8	46.5	78.8	28.5	71.1	46.1

TABLE IX
CATEGORICAL EMOTION CLASSIFICATION RESULTS ON THE **IEMOCAP** TEST SET (ACC. AND F1 INDICATE STANDARD ACCURACY AND F1 SCORES RESPECTIVELY, AND [†] INDICATES THE BASELINE MODELS' RESULTS FROM LI ET AL. [84])

Model	Angry		Excited		Frustrated		Happy		Neutral		Sad		Average	
	Acc.↑	F1↑	Acc.↑	F1↑	Acc.↑	F1↑	Acc.↑	F1↑	Acc.↑	F1↑	Acc.↑	F1↑	Acc.↑	F1↑
LSTM Pred. Fusion [†]	71.2	49.4	79.3	57.2	68.2	51.5	67.2	37.6	66.5	47.0	78.2	54.0	71.8	49.5
Transf. Pred. Fusion [†]	81.9	50.7	85.3	57.3	60.5	49.3	85.2	37.6	72.4	49.7	87.4	57.4	78.8	50.3
EmoEmbs [81]	65.9	48.9	73.5	58.3	68.5	52.0	69.6	38.3	73.6	48.7	80.8	53.0	72.0	49.8
MuT [99] [†]	77.9	60.7	76.9	58.0	72.4	57.0	80.0	46.8	74.9	53.7	83.5	65.4	77.6	56.9
AMOA [83] [†]	82.5	53.4	85.8	57.9	74.4	56.5	88.6	47.0	73.2	49.6	87.8	64.5	82.1	54.8
AVL Context Fusion	85.4	51.6	82.4	52.2	63.4	50.0	85.6	35.5	71.6	49.3	84.6	56.8	78.8	49.2
AVL COLD Fusion	86.5	58.4	88.3	61.6	68.6	57.9	88.1	43.7	76.6	48.9	87.9	60.2	82.7	55.1

the unimodal and multimodal emotion classification models. Note that this calibration error metric applies only to the classification settings. By computing the ECE values before and after applying temperature scaling to the softmax distributions over the predictions of each model separately, we analyse the impact of explicit uncertainty calibration (temperature scaling). We searched for an optimal temperature value in the range of $1e - 2$ to 1000 by doing a random search for 100 iterations. Similar to the technique followed in [14], we selected a temperature value that achieves the lowest ECE value on the validation set.

It is important to consider that the COLD fusion models are trained to be implicitly calibrated (see (6)) in terms of their context variance values. Thus, even before applying explicit calibration, i.e., temperature scaling, we expect the predictive uncertainty values or class-wise confidence scores of the COLD fusion models to have lower ECE values compared to the other fusion baselines.

Table V reports the ECE values for valence and arousal attributes on the AVEC 2019 corpus. For both attributes, before the application of temperature scaling, COLD fusion has the lowest calibration error when compared to the other models. After applying temperature scaling, it is obvious that the ECE values for all the models go down, and the COLD fusion still achieves the lowest error. Only in the case of valence, AV context fusion has a marginally lower ECE value compared to the COLD fusion. This minor discrepancy could be due to the random search of optimal temperature values and note that

here, different models have different optimal temperature values that are tuned for valence and arousal, separately. Nevertheless, in all the remaining cases (both before and after temperature scaling), COLD fusion consistently shows lower uncertainty calibration errors w.r.t. the other fusion models. Results on the IEMOCAP corpus (see Table VI) show similar trends, validating the effectiveness of the COLD fusion approach in producing well-calibrated uncertainty estimates. To visually illustrate the uncertainty calibration performance of the COLD fusion model, Appendix D, available online compares the reliability diagrams of different unimodal and multimodal dimensional emotion classification models.

Analysis of Audiovisual Fusion Weights: Fig. 6 illustrates modality-wise fusion weights estimated by the COLD fusion model on a validation sequence taken from the AVEC 2019 corpus. Note that these fusion weights are functions of the unimodal temporal context distributions (see (2)). In this illustration, we analyse the temporal patterns of fusion weights along with their corresponding unimodal and multimodal emotion predictions and their ground truth labels. This analysis clearly shows the well-calibrated nature of modality-wise fusion weights: when the predictions of one modality move closer to the ground truth compared to those of the other modality, the audiovisual weight values in the COLD fusion are found to be varying accordingly. From the transition points marked in Fig. 6, we can see that the fusion weights are gradually inverted, as the predictions of one modality move closer to the ground truth while the other modality predictions move further. This result validates our main




Fig. 6. Emotion predictions on an example from the AVEC 2019 CES validation set: Unimodal and multimodal valence predictions, and their uncertainty-based fusion weights estimated by the AV COLD fusion predictions. Note that fusion weights of the audio and visual modalities demonstrate (a) the calibration property – how far their corresponding unimodal predictions are from the ground truth ratings and (b) the ordinal ranking property – how well they can order the audio and visual modalities in terms of their reliability.

hypothesis of making unimodal latent distributions calibrated and ordinal for improved fusion performance.

VII. CONCLUSION

We proposed an uncertainty-aware multimodal fusion approach to dimensional and categorical emotion recognition from multimodal data. To capture modality-wise uncertainty w.r.t. predicting valence and arousal dimensions, we probabilistically modelled the unimodal temporal context by learning modality-wise latent distributions. For effective uncertainty-weighted multimodal fusion, we suggested conditioning the unimodal latent distributions such that their variance norms are learnt to be *well-calibrated* and *well-ranked (ordinal)*. To jointly impose these two constraints on the latent distributions, we introduced a novel softmax distributional matching loss function that encourages the uncertainty scores to be well-calibrated and well-ranked. Our novel loss function for multimodal learning is applicable to both classification and regression settings.

For example, in dimensional emotion regression tasks, COLD fusion shows $\sim 6\%$ average relative improvement over the best performing fusion baseline. Similarly, in the case of categorical

emotion classification, COLD fusion achieves $\sim 8.2\%$ relative accuracy improvement over the existing state-of-the-art model. Furthermore, we assess the robustness of different fusion models at test time by inducing noise into the visual modality through face masking. With the faces masked in 50% of the evaluation sequences, COLD fusion achieves $\sim 17\%$ average relative improvement over the best fusion baseline.

On spontaneous and acted emotion recognition tasks (in both dimensional and categorical emotion cases), our proposed uncertainty-aware fusion model achieved considerably better recognition performance than the uncertainty-unaware model-agnostic fusion baselines. In recognising dimensional emotions, COLD fusion demonstrated $\sim 6\%$ relative improvement over the best-performing fusion baseline, and in the case of categorical emotion recognition it achieved $\sim 8.2\%$ relative improvement over the existing state-of-the-art model. Validating our main hypothesis, extensive ablation studies (see Appendix B, available online) showed that it is important to apply both calibration and ordinality constraints for improving the emotion recognition results of uncertainty-aware fusion models. Furthermore, our method demonstrated noticeable improvements in terms of predictive uncertainty calibration errors of the emotion recognition models. It is important to note that our proposed calibration and ordinal ranking constraints can be easily applied to general model-fusion methods as well by quantifying the model-wise predictive uncertainty values of emotion labels. Future work can consider evaluating the COLD fusion approach on other complex multimodal learning tasks such as audiovisual speech recognition in noisy conditions [103] and humour detection [25], [104], etc.

In summary, this work showed the importance of uncertainty modelling for the dynamic integration of emotional expression cues from multimodal signals. We believe that uncertainty-aware information fusion is fundamental to reliably recognise apparent emotional states in naturalistic conditions. We hope that the results we demonstrated in this work may help in generating more interest in embracing uncertainty in multimodal affective computing.

REFERENCES

- [1] E. Hüllermeier and W. Waegeman, “Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods,” *Mach. Learn.*, vol. 110, pp. 457–506, 2021.
- [2] B. Schuller, M. Valster, F. Eyben, R. Cowie, and M. Pantic, “AVEC 2012: The continuous audio/visual emotion challenge,” in *Proc. 14th ACM Int. Conf. Multimodal Interact.*, 2012, pp. 449–456.
- [3] M. Valstar et al., “AVEC 2013: The continuous audio/visual emotion and depression recognition challenge,” in *Proc. AVEC Workshop*, 2013, pp. 3–10.
- [4] F. Ringeval et al., “AVEC 2019 workshop and challenge: State-of-mind, detecting depression with ai, and cross-cultural affect recognition,” in *Proc. AVEC Workshop*, 2019, pp. 3–12.
- [5] S. K. D’mello and J. Kory, “A review and meta-analysis of multimodal affect detection systems,” *ACM Comput. Surv.*, vol. 47, no. 3, pp. 1–36, 2015.
- [6] M. A. Nicolaou, H. Gunes, and M. Pantic, “Continuous prediction of spontaneous affect from multiple cues and modalities in valence-arousal space,” *IEEE Trans. Control Syst. Technol.*, vol. 2, no. 2, pp. 92–105, Apr.-Jun. 2011.
- [7] Z. Zeng, M. Pantic, G. I. Roisman, and T. S. Huang, “A survey of affect recognition methods: Audio, visual, and spontaneous expressions,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 1, pp. 39–58, Jan. 2009.

- [8] P. V. Rouast, M. T. Adam, and R. Chiong, "Deep learning for human affect recognition: Insights and new developments," *IEEE Trans. Control Syst. Technol.*, vol. 12, no. 2, pp. 524–543, Apr.-Jun. 2021.
- [9] F. Noroozi, M. Marjanovic, A. Njegus, S. Escalera, and G. Anbarjafari, "Audio-visual emotion recognition in video clips," *IEEE Trans. Control Syst. Technol.*, vol. 10, no. 1, pp. 60–75, Jan.-Mar. 2021.
- [10] L. Schoneveld, A. Othmani, and H. Abdelkawy, "Leveraging recent advances in deep learning for audio-visual emotion recognition," *Pattern Recognit. Lett.*, vol. 146, pp. 1–7, 2021.
- [11] M. Gerczuk, S. Amiriparian, S. Ottl, and B. W. Schuller, "EmoNet: A transfer learning framework for multi-corpus speech emotion recognition," *IEEE Trans. Control Syst. Technol.*, vol. 14, no. 2, pp. 1472–1487, Apr.-Jun. 2023.
- [12] O. Laurent et al., "Packed-ensembles for efficient uncertainty estimation," 2022, *arXiv:2210.09184*.
- [13] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 1321–1330.
- [14] J. Mukhoti, V. Kulharia, A. Sanyal, S. Golodetz, P. Torr, and P. Dokania, "Calibrating deep neural networks using focal loss," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 15288–15299.
- [15] A. Nguyen, J. Yosinski, and J. Clune, "Deep neural networks are easily fooled: High confidence predictions for unrecognizable images," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 427–436.
- [16] C. Szegedy et al., "Intriguing properties of neural networks," 2013, *arXiv:1312.6199*.
- [17] A. Kumar, S. Sarawagi, and U. Jain, "Trainable calibration measures for neural networks from kernel mean embeddings," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 2805–2814.
- [18] J. Moon, J. Kim, Y. Shin, and S. Hwang, "Confidence-aware learning for deep neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 7034–7044.
- [19] X. Yang, P. Ramesh, R. Chitta, S. Madhvanath, E. A. Bernal, and J. Luo, "Deep multimodal representation learning from temporal data," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 5447–5455.
- [20] C. Busso et al., "IEMOCAP: Interactive emotional dyadic motion capture database," *Lang. Resour. Eval.*, vol. 42, no. 4, pp. 335–359, 2008.
- [21] A. B. Zadeh, P. P. Liang, S. Poria, E. Cambria, and L.-P. Morency, "Multimodal language analysis in the wild: CMU-MOSEI dataset and interpretable dynamic fusion graph," in *Proc. Annu. Meeting Assoc. Comput. Linguistics*, 2018, pp. 2236–2246.
- [22] H. Gunes, B. Schuller, M. Pantic, and R. Cowie, "Emotion representation, analysis and synthesis in continuous space: A survey," in *Proc. IEEE Int. Conf. Autom. Face Gesture Recognit.*, 2011, pp. 827–834.
- [23] L. Stappen et al., "Muse 2020 challenge and workshop: Multimodal sentiment analysis, emotion-target engagement and trustworthiness detection in real-life media: Emotional car reviews in-the-wild," in *Proc. 1st Int. Multimodal Sentiment Anal. Real-life Media Challenge Workshop*, 2020, pp. 35–44.
- [24] L. Stappen et al., "The muse 2021 multimodal sentiment analysis challenge: Sentiment, emotion, physiological-emotion, and stress," in *Proc. 2nd Multimodal Sentiment Anal. Challenge*, 2021, pp. 5–14.
- [25] L. Christ et al., "The muse 2022 multimodal sentiment analysis challenge: Humor, emotional reactions, and stress," in *Proc. 3rd Int. Multimodal Sentiment Anal. Workshop Challenge*, 2022, pp. 5–14.
- [26] D. Kollias, A. Schulc, E. Hajiyev, and S. Zafeiriou, "Analysing affective behavior in the first ABAW 2020 competition," in *Proc. IEEE Int. Conf. Autom. Face Gesture Recognit.*, 2020, pp. 637–643.
- [27] D. Kollias and S. Zafeiriou, "Analysing affective behavior in the second ABAW2 competition," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 3652–3660.
- [28] T. Mittal, P. Guhan, U. Bhattacharya, R. Chandra, A. Bera, and D. Manocha, "Emoticon: Context-aware multimodal emotion recognition using Frege's principle," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 14234–14243.
- [29] D. Yang et al., "Emotion recognition for multiple context awareness," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 144–162.
- [30] S. Poria, E. Cambria, R. Bajpai, and A. Hussain, "A review of affective computing: From unimodal analysis to multimodal fusion," *Inf. Fusion*, vol. 37, pp. 98–125, 2017.
- [31] Y. Jiang, W. Li, M. S. Hossain, M. Chen, A. Alelaiwi, and M. Al-Hammadi, "A snapshot research and implementation of multimodal information fusion for data-driven emotion recognition," *Inf. Fusion*, vol. 53, pp. 209–221, 2020.
- [32] S. Zhao, G. Jia, J. Yang, G. Ding, and K. Keutzer, "Emotion recognition from multiple modalities: Fundamentals and methodologies," *IEEE Signal Process. Mag.*, vol. 38, no. 6, pp. 59–73, Nov. 2021.
- [33] J. She, Y. Hu, H. Shi, J. Wang, Q. Shen, and T. Mei, "Dive into ambiguity: Latent distribution mining and pairwise uncertainty estimation for facial expression recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 6248–6257.
- [34] Y. Zhang, C. Wang, and W. Deng, "Relative uncertainty learning for facial expression recognition," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2021, pp. 17616–17627.
- [35] K. Wang, X. Peng, J. Yang, S. Lu, and Y. Qiao, "Suppressing uncertainties for large-scale facial expression recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 6897–6906.
- [36] N. Foteinopoulou, C. Tzelepis, and I. Patras, "Estimating continuous affect with label uncertainty," in *Proc. Int. Conf. Affect. Comput. Intell. Interact.*, 2021, pp. 1–8.
- [37] Z. Zeng, J. Tu, M. Liu, and T. S. Huang, "Multi-stream confidence analysis for audio-visual affect recognition," in *Proc. Int. Conf. Affect. Comput. Intell. Interact.*, 2005, pp. 964–971.
- [38] Z. Xie and L. Guan, "Multimodal information fusion of audiovisual emotion recognition using novel information theoretic tools," in *Proc. IEEE Int. Conf. Multimedia Expo*, 2013, pp. 1–6.
- [39] T. Dang et al., "Investigating word affect features and fusion of probabilistic predictions incorporating uncertainty in AVEC 2017," in *Proc. AVEC workshop*, 2017, pp. 27–35.
- [40] E. Sanchez, M. K. Tellamekala, M. Valstar, and G. Tzimiropoulos, "Affective processes: Stochastic modelling of temporal context for emotion and facial expression recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 9070–9080.
- [41] T. Mani Kumar, E. Sanchez, G. Tzimiropoulos, T. Giesbrecht, and M. Valstar, "Stochastic process regression for cross-cultural speech emotion recognition," in *Proc. Interspeech*, 2021, pp. 3390–3394.
- [42] M. Garnelo et al., "Conditional neural processes," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 1704–1713.
- [43] M. Garnelo et al., "Neural processes," 2018, *arXiv:1807.01622*.
- [44] M. K. Tellamekala, T. Giesbrecht, and M. Valstar, "Modelling stochastic context of audio-visual expressive behaviour with affective processes," *IEEE Trans. Control Syst. Technol.*, vol. 14, no. 3, pp. 2290–2303, Jul.-Sep. 2023.
- [45] A. Schörgendorfer and W. Elmenreich, "Extended confidence-weighted averaging in sensor fusion," in *Proc. Junior Scientist Conf.*, 2006, pp. 67–68.
- [46] P. W. Große, H. Holzapfel, and A. Waibel, "Confidence based multimodal fusion for person identification," in *Proc. ACM Multimedia*, 2008, pp. 885–888.
- [47] G. Papandreou, A. Katsamanis, V. Pitsikalis, and P. Maragos, "Adaptive multimodal fusion by uncertainty compensation with application to audiovisual speech recognition," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 17, no. 3, pp. 423–435, Mar. 2009.
- [48] M. Subedar, R. Krishnan, P. L. Meyer, O. Tickoo, and J. Huang, "Uncertainty-aware audiovisual activity recognition using deep Bayesian variational inference," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 6301–6310.
- [49] J. Tian, W. Cheung, N. Glaser, Y.-C. Liu, and Z. Kira, "UNO: Uncertainty-aware noisy-or multimodal fusion for unanticipated input degradation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 5716–5723.
- [50] S. Pramanick, A. Roy, and V. M. Patel, "Multimodal learning using optimal transport for sarcasm and humor detection," in *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, 2022, pp. 3930–3940.
- [51] H. Wang et al., "Uncertainty-aware multi-modal learning via cross-modal random network prediction," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 200–217.
- [52] U. Sarawagi, R. Khincha, W. Zulfikar, S. Ghosh, and P. Maes, "Uncertainty-aware boosted ensembling in multi-modal settings," in *Proc. Int. Joint Conf. Neural Netw.*, 2021, pp. 1–9.
- [53] T. Zhi-Xuan, H. Soh, and D. Ong, "Factorized inference in deep Markov models for incomplete multimodal time series," in *Proc. AAAI Conf. Artif. Intell.*, 2020, pp. 10334–10341.
- [54] D.-B. Wang, L. Feng, and M.-L. Zhang, "Rethinking calibration of deep neural networks: Do not be afraid of overconfidence," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2021, pp. 11809–11820.
- [55] M. Minderer et al., "Revisiting the calibration of modern neural networks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2021, pp. 15682–15694.

- [56] B. Zadrozny and C. Elkan, "Obtaining calibrated probability estimates from decision trees and naive Bayesian classifiers," in *Proc. Int. Conf. Mach. Learn.*, 2001, pp. 609–616.
- [57] B. Zadrozny and C. Elkan, "Transforming classifier scores into accurate multiclass probability estimates," in *Proc. 8th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2002, pp. 694–699.
- [58] J. Platt et al., "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods," in *Proc. Adv. Large Margin Classifiers*, 1999, pp. 61–74.
- [59] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," 2015, *arXiv:1503.02531*.
- [60] R. Krishnan and O. Tickoo, "Improving model calibration with accuracy versus uncertainty optimization," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 18237–18248.
- [61] V. Kuleshov, N. Fenner, and S. Ermon, "Accurate uncertainties for deep learning using calibrated regression," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 2796–2804.
- [62] H. Song, T. Diethe, M. Kull, and P. Flach, "Distribution calibration for regression," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 5897–5906.
- [63] S. Utpala and P. Rai, "Quantile regularization: Towards implicit calibration of regression models," 2020, *arXiv:2002.12860*.
- [64] W. Li, X. Huang, J. Lu, J. Feng, and J. Zhou, "Learning probabilistic ordinal embeddings for uncertainty-aware regression," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 13896–13905.
- [65] C. Corbière, N. Thome, A. Bar-Hen, M. Cord, and P. Pérez, "Addressing failure prediction by learning model confidence," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 2898–2909.
- [66] R. Roady, T. L. Hayes, R. Kemker, A. Gonzales, and C. Kanan, "Are out-of-distribution detection methods effective on large-scale datasets?" 2019, *arXiv:1910.14034*.
- [67] Y. Geifman and R. El-Yaniv, "Selective classification for deep neural networks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 4878–4887.
- [68] T. Baltrušaitis, C. Ahuja, and L.-P. Morency, "Multimodal machine learning: A survey and taxonomy," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 2, pp. 423–443, Feb. 2019.
- [69] S. Zhang, S. Zhang, T. Huang, W. Gao, and Q. Tian, "Learning affective features with a hybrid deep model for audio-visual emotion recognition," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 28, no. 10, pp. 3030–3043, Oct. 2018.
- [70] F. Lingenfelser, J. Wagner, J. Deng, R. Brueckner, B. Schuller, and E. André, "Asynchronous and event-based fusion systems for affect recognition on naturalistic data in comparison to conventional approaches," *IEEE Trans. Control Syst. Technol.*, vol. 9, no. 4, pp. 410–423, Oct.-Dec. 2018.
- [71] F. Ringeval et al., "Prediction of asynchronous dimensional emotion ratings from audiovisual and physiological data," *Pattern Recognit. Lett.*, vol. 66, pp. 22–30, 2015.
- [72] P. Tzirakis, G. Trigeorgis, M. A. Nicolaou, B. W. Schuller, and S. Zafeiriou, "End-to-end multimodal emotion recognition using deep neural networks," *IEEE J. Selected Topics Signal Process.*, vol. 11, no. 8, pp. 1301–1309, Dec. 2017.
- [73] G. Evangelopoulos et al., "Multimodal saliency and fusion for movie summarization based on aural, visual, and textual attention," *IEEE Trans. Multimedia*, vol. 15, no. 7, pp. 1553–1568, Nov. 2013.
- [74] J. Kossaifi, A. Toisoul, A. Bulat, Y. Panagakis, T. M. Hospedales, and M. Pantic, "Factorized higher-order CNNs with an application to spatio-temporal emotion estimation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 6059–6068.
- [75] S. Bruch, X. Wang, M. Bendersky, and M. Najork, "An analysis of the softmax cross entropy loss for learning-to-rank with binary relevance," in *Proc. ACM SIGIR Int. Conf. Theory Inf. Retrieval*, 2019, pp. 75–78.
- [76] J. Chang, Z. Lan, C. Cheng, and Y. Wei, "Data uncertainty learning in face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 5710–5719.
- [77] J. Kossaifi et al., "SEWA DB: A rich database for audio-visual emotion and sentiment research in the wild," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 3, pp. 1022–1040, Mar. 2021.
- [78] Z. Zhao, Y. Zheng, Z. Zhang, H. Wang, Y. Zhao, and C. Li, "Exploring spatio-temporal representations by integrating attention-based bidirectional-LSTM-RNNS and FCNS for speech emotion recognition," in *Proc. Interspeech*, 2018, pp. 272–276.
- [79] S. Tripathi and H. Beigi, "Multi-modal emotion recognition on IEMO-CAP with neural networks," 2018, *arXiv:1804.05788*.
- [80] W. Dai, S. Cahyawijaya, Z. Liu, and P. Fung, "Multimodal end-to-end sparse model for emotion recognition," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics*, 2021, pp. 5305–5316.
- [81] W. Dai, Z. Liu, T. Yu, and P. Fung, "Modality-transferable emotion embeddings for low-resource multimodal emotion recognition," in *Proc. Conf. Empir. Methods Natural Lang. Joint Conf. Natural Lang. Process.*, 2020, pp. 269–280.
- [82] I. Lawrence and K. Lin, "A concordance correlation coefficient to evaluate reproducibility," *Biometrics*, vol. 45, pp. 255–268, 1989.
- [83] Z. Li et al., "AMOA: Global acoustic feature enhanced modal-order-aware network for multimodal sentiment analysis," in *Proc. 29th Int. Conf. Comput. Linguistics*, 2022, pp. 7136–7146.
- [84] Z. Li, Y. Zhou, Y. Liu, F. Zhu, C. Yang, and S. Hu, "QAP: A quantum-inspired adaptive-priority-learning model for multimodal emotion recognition," in *Proc. Annu. Meeting Assoc. Comput. Linguistics*, 2023, pp. 12191–12204.
- [85] J. Yang, A. Bulat, and G. Tzimiropoulos, "FAN-face: A simple orthogonal improvement to deep face recognition," in *Proc. AAAI Conf. Artif. Intell.*, 2020, pp. 12621–12628.
- [86] A. Toisoul, J. Kossaifi, A. Bulat, G. Tzimiropoulos, and M. Pantic, "Estimation of continuous valence and arousal levels from faces in naturalistic conditions," *Nature Mach. Intell.*, vol. 3, no. 1, pp. 42–50, 2021.
- [87] I. Ntinou, E. Sanchez, A. Bulat, M. Valstar, and Y. Tzimiropoulos, "A transfer learning approach to heatmap regression for action unit intensity estimation," *IEEE Trans. Control Syst. Technol.*, vol. 14, no. 1, pp. 436–450, Jan.-Mar. 2023.
- [88] A. Mollahosseini, B. Hasani, and M. H. Mahoor, "AffectNet: A database for facial expression, valence, and arousal computing in the wild," *IEEE Trans. Control Syst. Technol.*, vol. 10, no. 1, pp. 18–31, Jan.-Mar. 2019.
- [89] H. Chen, Y. Deng, S. Cheng, Y. Wang, D. Jiang, and H. Sahli, "Efficient spatial temporal convolutional features for audiovisual continuous affect recognition," in *Proc. AVEC Workshop*, 2019, pp. 19–26.
- [90] S. Hershey et al., "CNN architectures for large-scale audio classification," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2017, pp. 131–135.
- [91] S. Chen, Q. Jin, J. Zhao, and S. Wang, "Multimodal multi-task learning for dimensional and continuous emotion recognition," in *Proc. AVEC Workshop*, 2017, pp. 19–26.
- [92] D. S. Park et al., "SpecAugment: A simple data augmentation method for automatic speech recognition," in *Proc. Interspeech*, 2019, pp. 2613–2617.
- [93] G. Degottex, J. Kane, T. Drugman, T. Raitio, and S. Scherer, "COVAREP—a collaborative voice analysis repository for speech technologies," in *Proc. Int. Conf. Acoust. Speech Signal Process.*, 2014, pp. 960–964.
- [94] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Proc. Conf. Empir. Methods Natural Lang. Joint Conf. Natural Lang. Process.*, 2014, pp. 1532–1543.
- [95] I. Loshchilov and F. Hutter, "SGDR: Stochastic gradient descent with warm restarts," 2016, *arXiv:1608.03983*.
- [96] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [97] R. Liaw, E. Liang, R. Nishihara, P. Moritz, J. E. Gonzalez, and I. Stoica, "Tune: A research platform for distributed model selection and training," 2018, *arXiv:1807.05118*.
- [98] J. Zhao, R. Li, J. Liang, S. Chen, and Q. Jin, "Adversarial domain adaption for multi-cultural dimensional emotion recognition in dyadic interactions," in *Proc. AVEC Workshop*, 2019, pp. 37–45.
- [99] 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 *Proc. Annu. Meeting Assoc. Comput. Linguistics*, 2019, Art. no. 6558.
- [100] A. Anwar and A. Raychowdhury, "Masked face recognition for secure authentication," 2020, *arXiv:2008.11104*.
- [101] Q. Wei, X. Huang, and Y. Zhang, "FV2ES: A fully end2end multimodal system for fast yet effective video emotion recognition inference," *IEEE Trans. Broadcast.*, vol. 69, no. 1, pp. 10–20, Mar. 2023.
- [102] D. Hazarika, R. Zimmermann, and S. Poria, "MISA: Modality-invariant and-specific representations for multimodal sentiment analysis," in *Proc. ACM Multimedia*, 2020, pp. 1122–1131.
- [103] D. Stewart, R. Seymour, A. Pass, and J. Ming, "Robust audio-visual speech recognition under noisy audio-video conditions," *IEEE Trans. Cybern.*, vol. 44, no. 2, pp. 175–184, Feb. 2014.
- [104] L. Christ et al., "The muse 2023 multimodal sentiment analysis challenge: Mimicked emotions, cross-cultural humour, and personalisation," 2023, *arXiv:2305.03369*.
- [105] Y. Wang, Y. Shen, Z. Liu, P. P. Liang, A. Zadeh, and L.-P. Morency, "Words can shift: Dynamically adjusting word representations using non-verbal behaviors," in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 7216–7223.

- [106] H. Pham, P. P. Liang, T. Manzini, L.-P. Morency, and B. Póczos, “Found in translation: Learning robust joint representations by cyclic translations between modalities,” in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 6892–6899.
- [107] Z. Liu, Y. Shen, V. B. Lakshminarasimhan, P. P. Liang, A. Zadeh, and L.-P. Morency, “Efficient low-rank multimodal fusion with modality-specific factors,” in *Proc. Annu. Meeting Assoc. Comput. Linguistics*, 2018, pp. 2247–2256.
- [108] A. Zadeh, M. Chen, S. Poria, E. Cambria, and L.-P. Morency, “Tensor fusion network for multimodal sentiment analysis,” in *Proc. Conf. Empir. Methods Natural Lang. Joint Conf. Natural Lang. Process.*, 2017, pp. 1103–1114.
- [109] Y.-H. H. Tsai, P. P. Liang, A. Zadeh, L.-P. Morency, and R. Salakhutdinov, “Learning factorized multimodal representations,” 2018, *arXiv:1806.06176*.
- [110] Z. Sun, P. Sarma, W. Sethares, and Y. Liang, “Learning relationships between text, audio, and video via deep canonical correlation for multimodal language analysis,” in *Proc. AAAI Conf. Artif. Intell.*, 2020, pp. 8992–8999.
- [111] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics Hum. Lang. Technol.*, 2019, pp. 4171–4186.



Mani Kumar Tellamekala received the bachelor’s degree in electronics and communication engineering from IIIT-RGUKT RK Valley, India. He is currently working toward the PhD degree with the Computer Vision Laboratory, University of Nottingham, U.K.. His research focuses on uncertainty-aware temporal learning and self-supervised representation learning, and their applications to affect recognition from face and voice data.



Shahin Amiriparian (Member, IEEE) received the doctoral degree with the highest honours (summa cum laude) from the Technical University of Munich, Germany, in 2019. Currently, he is a postdoctoral researcher with the Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, Germany. His main research focus is deep learning, unsupervised representation learning, and transfer learning for machine perception, affective computing, and audio understanding.



Björn W. Schuller (Fellow, IEEE) received the diploma, in 1999, the doctoral degree, in 2006, and the Habilitation and Adjunct Teaching Professorship in the subject area of signal processing and machine intelligence, in 2012, all in electrical engineering and information technology from the Technical University of Munich, Germany. He is professor of Artificial Intelligence in the Department of Computing, Imperial College London, U.K., a full professor and head of the chair of Embedded Intelligence for Health Care and Wellbeing at the University of Augsburg, Germany. He (co-)authored five books and more than 1200 publications in peer-reviewed books, journals, and conference proceedings leading to more than 46000 citations (h-index = 97).



Elisabeth André is a full professor of computer science and founding chair of human-centered artificial intelligence with Augsburg University, Augsburg, Germany. Her work has won many awards including Gottfried Wilhelm Leibnitz Prize 2021, with 2.5 Mio € the highest endowed German research award. In 2010, she was elected a member of the Academy of Europe and the German Academy of Sciences Leopoldina. In 2017, she was elected to the CHI Academy, an honorary group of leaders in the field of human-computer interaction. To honor her achievements in bringing artificial intelligence techniques to human-computer interaction, she was awarded a EurAI fellowship, in 2013. From 2019 - 2022, she has been serving as the editor-in-chief of *IEEE Transactions on Affective Computing*.



Timo Giesbrecht received the PhD degree from Maastricht University. He is a consumer science leader in Beauty & Personal Care Science & Technology with Unilever R&D Port Sunlight. In his work, he employs a wide array of measures and paradigms such as sensory, neuroimaging, psychophysiological measures, and behavioural tasks in collaboration with leading academics around the world. He published more than 95 scientific papers and book chapters and his Hirsch Index is 35. For his research, he has received funding from Dutch and U.K. Research Councils. Before joining Unilever, Timo Giesbrecht worked with Maastricht University, the Netherlands and Mount Sinai School of Medicine, New York.



Michel Valstar is professor in Automatic Understanding of Human Behaviour with the University of Nottingham, School of Computer Science. He is a member of both the Computer Vision Lab and the Mixed Reality Lab. Automatic Human Behaviour Understanding encompasses Machine Learning analysis of data from observable human behaviour and a good idea of how and why people behave in this world. Valstar was a visiting Researcher with the Affective Computing group, the Media Lab, MIT, and a research associate with the iBUG group, Department of Computing, Imperial College London. His expertise is facial expression recognition, in particular the analysis of FACS Action Units. He recently proposed a new field of research called ‘Behaviomedics’, which applies affective computing and Social Signal Processing to the field of medicine to help diagnose, monitor, and treat medical conditions that alter expressive behaviour such as depression.