

RESEARCH ARTICLE

emoDARTS: Joint Optimization of CNN and Sequential Neural Network Architectures for Superior Speech Emotion Recognition

THEJAN RAJAPAKSHE¹, (Student Member, IEEE), **RAJIB RANA¹**, (Member, IEEE),
SARA KHALIFA², (Member, IEEE), **BERRAK SISMAN³**, (Member, IEEE),
BJÖRN W. SCHULLER^{4,5}, (Fellow, IEEE), AND **CARLOS BUSO⁶**, (Fellow, IEEE)

¹School of Mathematics, Physics and Computing, University of Southern Queensland, Toowoomba, QLD 4350, Australia

²School of Information Systems, Queensland University of Technology, Brisbane, QLD 4000, Australia

³Speech & Machine Learning (SML) Lab, The University of Texas at Dallas, Richardson, TX 75080, USA

⁴Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, 86159 Augsburg, Germany

⁵GLAM—Group on Language, Audio and Music, Imperial College London, London SW7 2AZ London, U.K.

⁶Department of Electrical and Computer Engineering, The University of Texas at Dallas, Richardson, TX 75080, USA

Corresponding author: Thejan Rajapakshe (thejan.rajapakshe@unisoq.edu.au)

ABSTRACT Speech Emotion Recognition (SER) is crucial for enabling computers to understand the emotions conveyed in human communication. With recent advancements in Deep Learning (DL), the performance of SER models has significantly improved. However, designing an optimal DL architecture requires specialised knowledge and experimental assessments. Fortunately, Neural Architecture Search (NAS) provides a potential solution for automatically determining the best DL model. The Differentiable Architecture Search (DARTS) is a particularly efficient method for discovering optimal models. This study presents emoDARTS, a DARTS-optimised joint CNN and Sequential Neural Network (SeqNN: LSTM, RNN) architecture that enhances SER performance. The literature supports the selection of CNN and LSTM coupling to improve performance. While DARTS has previously been used to choose CNN and LSTM operations independently, our technique adds a novel mechanism for selecting CNN and SeqNN operations in conjunction using DARTS. Unlike earlier work, we do not impose limits on the layer order of the CNN. Instead, we let DARTS choose the best layer order inside the DARTS cell. We demonstrate that emoDARTS outperforms conventionally designed CNN-LSTM models and surpasses the best-reported SER results achieved through DARTS on CNN-LSTM by evaluating our approach on the IEMOCAP, MSP-IMPROV, and MSP-Podcast datasets.

INDEX TERMS Speech emotion recognition, neural architecture search, deep learning, DARTS.

I. INTRODUCTION

Recognising the emotional nuances embedded in speech is a fundamental, yet complex challenge. Over the last decade, the field of Speech Emotion Recognition (SER) has experienced significant strides, predominantly driven by the exponential growth of deep learning [1], [2], [3], [4]. A key breakthrough facilitated by deep learning is its capability to automatically learn features, departing from the traditional reliance on manually crafted features shaped by human perceptions

The associate editor coordinating the review of this manuscript and approving it for publication was Cesar Vargas-Rosales¹.

of speech signals. Nevertheless, determining the optimal deep-learning architecture for SER remains a challenging task that warrants attention. Conventional approaches involve iterative modifications and recursive training of models until an optimal configuration is found. However, this approach becomes prohibitively time-consuming due to the extensive training and testing required for numerous configurations.

An alternative to the conventional approach is the “Neural Architecture Search” (NAS), which can help discover optimal neural networks for a given task. The idea is to find the models’ architecture to minimise the loss. In NAS, search is done over a discrete set of candidate operations, which




FIGURE 1. The proposed architecture of emoDARTS passes the input features to the CNN component through the SeqNN component and finally to a dense layer. The optimum CNN and SeqNN operations are selected by DARTS jointly.

requires the model to be trained on a specific configuration before moving on to the next configuration. Nevertheless, this approach demands considerable time and resources.

The Differentiable Architecture Search (DARTS) is a method that has been developed to optimise the search for a neural network architecture. It allows for the relaxation of the discrete set of candidate operations, making the space continuous and reducing the computation time significantly, from 2,000 GPU days to just 2-3 GPU days. This is a major improvement from the previous methods of reinforcement learning or evolution algorithm, which required 2,000 and 3,150 GPU days, respectively. Additionally, through network optimisation DARTS has the potential to offer significantly high SER accuracy, which is currently quite low and needs improvement. These two points serve as motivation to use DARTS for SER.

Additionally, previous studies have shown that a multi-temporal Convolutional Neural Network (CNN) stacked on a Long Short-Term Memory Network (LSTM) can capture contextual information at multiple temporal resolutions, complementing LSTM for modelling long-term contextual information, thus offering improved performance [5], [6], [7], [8]. Sequential Neural Networks (SeqNN) like Recurrent Neural Networks (RNN) or LSTM can easily identify the patterns of a sequential stream of data. This paper takes a pioneering step by leveraging DARTS for a novel joint CNN-SeqNN configuration, named “emoDARTS”, as depicted in Figure 1, with an attention network seamlessly integrated into the SeqNN component to further elevate its performance.

The investigation of DARTS within the SER domain is minimal and invites further inquiry to uncover the potential for improving SER performance. DARTS has only recently been employed in SER tasks to improve models, as recently as 2022 [9], [10], wherein the researchers have mostly applied DARTS separately on CNNs and RNNs [11], [12], [13]. Hence, the viability of utilising DARTS jointly for CNN and SeqNN requires exploration. While there is a lone study that explores the joint optimisation of CNN and LSTM [10], it imposes constraints on the layer order for the CNN within the DARTS component, thereby limiting the full potential of DARTS. In response to this limitation, our paper takes on the challenge of optimising this joint configuration without such

constraints. The contributions of this paper are summarised as follows.

- 1) This paper proposes a novel DARTS-optimised joint CNN and SeqNN architecture, emoDARTS, achieving greater autonomy to DARTS in selecting optimal network configurations.
- 2) We demonstrate the robust generalisation capabilities of the proposed emoDARTS model by testing it on three widely used datasets: IEMOCAP, MSP-IMPROV, and MSP-Podcast.
- 3) Experimental results demonstrate that emoDARTS achieves considerably higher SER accuracy than humans designing the CNN-LSTM configuration. It also outperforms the best-reported SER results achieved using DARTS on CNN-LSTM.

II. RELATED WORK

This section delves into the existing literature on using DARTS and NAS for SER. Notably, our exploration reveals a limited number of papers in this space. We therefore extend our review to encompass relevant papers in related fields to provide a comprehensive perspective. For completeness, we also include studies employing CNNs, LSTM networks, and their joint utilisation for SER.

A. SPEECH EMOTION RECOGNITION USING CNN AND LSTM

One of the earliest uses of CNN networks in SER is reported by Zheng et al. in 2015 [14]. The authors introduced a spectrum generated from an audio signal to a CNN network and output the recognised emotion. The authors report that they can surpass the SVM-based classification performance and reach 40% classification accuracy for a five-class classification using the IEMOCAP dataset.

The earliest work combining CNN and LSTM for SER is by Trigeorgis et al. in 2016 [15]. The authors show an impressive improvement by a fully self-learned representation over traditional expert-crafted features on dimensional emotion recognition.

Zhao et al. [1] show that using CNN and LSTM networks combined in the same SER model produces better results than using only CNN. Using the IEMOCAP dataset, they obtained a speaker-independent accuracy of 52% by using a log-Mel spectrogram as the input feature. Their SER approach utilises an LSTM layer to learn contextual dependencies in local features, while a CNN-based layer learns the local features.

B. APPLICATION OF NAS AND DARTS IN SER AND RELATED FIELDS

The first paper suggesting NAS in SER was by Zhang et al. in 2018 [16]. The authors employ a controller network that shapes the architecture by the number of layers and nodes per layer and the hyperparameter activation function of a child network by reinforcement learning. They show an

TABLE 1. Summary and focus on the literature on NAS, DARTS, and speech emotion recognition.

Paper	Focus				Dataset		
	SER	NAS	DARTS	Joint Opt. of CNN & SeqNN.	IEMOCAP	MSP-IMPROV	MSP-Podcast
Zoph and Le 2016 [17]	✗	✓	✗	✗	✗	✗	✗
Zhang et al. 2018 [14]	✓	✓	✗	✗	✓	✗	✗
Liu et al. 2018 [20]	✗	✓	✓	✗	✗	✗	✗
Gong et al. 2019 [19]	✗	✓	✗	✗	✗	✗	✗
Liberis et al. 2020 [18]	✗	✓	✗	✗	✗	✗	✗
Wu et al. 2022 [10]	✓	✓	✓	✓	✓	✗	✗
Sun et al. 2023 [9]	✓	✓	✓	✗	✓	✗	✗
emoDARTS	✓	✓	✓	✓	✓	✓	✓

improvement over human-designed architectures and random searches of these.

Zoph and Le [17] use reinforcement learning to optimise an RNN network that develops model architectures to maximise the resulting accuracy of the generated model. As a result, they develop outstanding models for the CIFAR-10 and Penn Treebank datasets. They were able to develop a convolutional network architecture for the CIFAR-10 dataset which has a 3.65 error rate and a recurrent network architecture for Penn Treebank with 62.4 perplexity.

Even though NAS is primarily used to find optimised architecture for complex and large models, researchers have also studied the possibility of using NAS to design smaller deep neural network models. Liberis et al. [18] develop a NAS system called μ NAS to design smaller neural architectures that can run on microcontroller units. They improve the top-1 accuracy by 4.8% in image classification problems while reducing the memory footprint up to 13 times. Similarly, Gong et al. [19] study the feasibility of using NAS for reducing deep learning models to deploy on resource-constrained edge devices.

Traditional NAS consumes much computational power and time to achieve the optimal model for a given problem. In 2018, Liu et al. [20] came up with a differentiable approach to solving the optimisation by continuous relaxation of the architecture representation. This approach is more compute efficient and high performing as the search space is not discrete and non-differentiable. They produce high-performing CNN and RNN architectures for tasks such as image recognition and language modelling within a fraction of the search cost of traditional NAS algorithms. DARTS has been popular in the past three years with many studies carried on for extending and improving the algorithm [11], [21], [22], [23], [24], [25].

Wu et al. [10] proposed a uniform path dropout strategy to optimise candidate architecture. They use SER as their DARTS application and the IEMOCAP dataset to develop an SER model with an accuracy of 56.28% for a four-class classification problem using discrete Fourier transform spectrograms extracted from audio as input. In their work, the authors specify layer order as two convolution layers at first,

followed by a max-pooling layer, a convolution layer. They use DARTS to select the optimum parameters for each layer. We, on the other hand, do not specify the layer sequence and instead enable DARTS to select the ideal design with minimal interference.

EmotionNAS is a two-branch NAS strategy introduced by Sun et al. [9] in 2023. The authors use DARTS to optimise their two models in two branches, the CNN model and RNN model, which use a spectrogram and a waveform as inputs, respectively. They obtained an unweighted accuracy of 72.1% from the combined model for the IEMOCAP dataset. They also report the performance of 63.2% in the spectrogram branch, which only uses a CNN component. The main difference between our approach and the study by Sun et al. [9] is that we use a SeqNN component coupled in series with the CNN layer as in Figure 2 while Sun et al. [9] use an RNN layer in parallel to the CNN layer in a different branch.

We conducted preliminary research to determine the feasibility of utilising DARTS for SER in a CNN-LSTM architecture, where we only optimised the CNN network using DARTS [26]. This paper extends the idea of using DARTS in SER but with more relaxation in the SeqNN component by jointly optimising the whole architecture.

In recent years, the literature has highlighted the use of attention networks in SER, which has provided superior outcomes [27], [28]. We added an attention network component to the DARTS search scope to discover whether it improves performance. Zou et al. [29] have introduced a concept called ‘co-attention’ where many separate inputs from multimodal inputs are fused by co-attention. They used three sets of features MFCC, spectrogram, and Wav2Vec2 features from the IEMOCAP dataset and obtained 72.70% accuracy. Liu et al. [30] have utilised an attention-based bi-directional LSTM followed by a CNN layer for a SER problem. They have achieved a significant performance of 66.27% for the IEMOCAP Dataset. Their idea of ‘CNN - LSTM attention’ paved the foundation for our model architecture.

In Table 1, we briefly compare the existing studies with emoDARTS. The comparison clearly shows that,

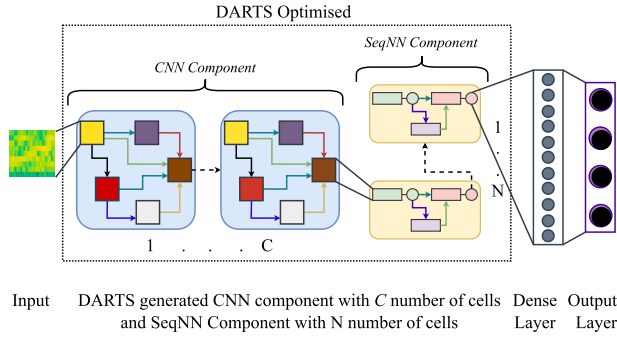


FIGURE 2. The emoDARTS architecture comprises input features processed through CNN, SeqNN, and Dense layers and it utilises DARTS for jointly optimising the CNN and SeqNN components.

- 1) While some studies employ NAS for SER, the utilisation of DARTS in SER is notably limited.
- 2) Singularly, one study has explored the concept of jointly optimising CNN and SeqNN using DARTS for SER, in which the researchers specified the layer order. However, in our study, we let DARTS determine the optimal network from a relaxed search scope which enables it to select any operation in search space at the optimum layer.
- 3) Most existing studies primarily focus on the IEMOCAP dataset. In contrast, our study uniquely incorporates three widely recognised SER datasets: IEMOCAP, MSP-IMPROV, and MSP-Podcast to demonstrate the generalisation power of emoDARTS.

III. emoDARTS FRAMEWORK

The proposed ‘emoDARTS’ uses DARTS to improve SER using a CNN-SeqNN network, which was motivated by studies that showed increased SER performance when CNN and LSTM layers were combined [5], [7], [8], [30]. We represent our network as a multi-component DARTS network, with the input fed into a CNN component and the output features from the CNN component fed into a SeqNN component, but all components are optimised jointly during the architecture search phase under a single compute graph, delivering an optimal architecture (Figure 2).

DARTS uses a differentiable approach to network optimisation. A computation cell is the DARTS algorithm’s fundamental unit. It aims to optimise the cell so that the architecture can function to its maximum performance. A DARTS cell is described as a directed graph, with each node representing a feature (representation) and each edge representing an operation that can be performed to a representation. One unique feature of this network is that each node is connected to all of its previous nodes by an edge, as seen in Figure 3 (a). If the output of the node j is $x^{(j)}$ and the operation ‘ o ’ on the edge connecting the nodes i and j is $o^{(i,j)}$, $x^{(j)}$ can be obtained by the Equation 1:

$$x^{(j)} = \sum_{i < j} o^{(i,j)}(x^{(i)}) \quad (1)$$

In the beginning, the candidate search space is generated by combining each node of the DARTS cell with all the candidate operations (with multiple links between nodes), as illustrated in Figure 3 (b). Equation 1 incorporates a weight parameter α to identify the optimal edge (operation) connecting two nodes, i and j , from the candidate search space of all operations. Equation 2 describes how the node’s output be represented.

$$x^{(j)} = \sum_{i < j} \alpha^{(i,j)} o^{(i,j)}(x^{(i)}) \quad (2)$$

Then, the continuous relaxation of the search space updates the weights ($\alpha^{(i,j)}$) of the edges. The final architecture can be obtained by selecting the operation between two nodes with the highest weight ($o^{(i,j)*}$) by using Equation 3.

$$o^{(i,j)*} = \operatorname{argmax}_o(\alpha^{(i,j)}) \quad (3)$$

The searched discrete cell architecture is shown in Figure 3 (d).

The number of cells (C or N) in a component is a parameter for the DARTS algorithm that specifies how many DARTS cells are stacked to form a component in the model. Each cell takes the last two cells’ output as input. If the output from each cell t is y_t and the function within the cell is f , then y_t can be represented as;

$$y_t = f(y_{t-1}, y_{t-2}) \quad (4)$$

DARTS’ CNN component has two types of CNN cells: ‘normal’ and ‘reduction’ cells. It sets the stride to one in normal cells and two in reduction cells, resulting in a down-sampled output in the reduction cells. This downsampling allows the model to eliminate the duplication of intermediate characteristics, reducing complexity.

We decided to jointly optimise the CNN and SeqNN components rather than individually since it is important for downstream components (in this case, the CNN component) to understand the behaviour of upstream components (SeqNN). Joint optimisation in a multiple-component network improves architecture search in various ways, including: 1. the back-propagation of loss minimisation flows through all the components in a single compute graph; and 2. it reduces the time required in the search phase by searching the architecture of the whole network at once.

IV. EXPERIMENTAL SETUP

A. DATASET AND FEATURE SELECTION

We use the widely used IEMOCAP [31], MSP-IMPROV [32], and MSP-Podcast [33] datasets for our experiments. Our study takes the improvised subset of IEMOCAP and the four categorical labels, happiness, sadness, anger, and neutral as classes from the datasets. We employ five-fold cross-validation with at least one speaker out in our training and evaluations. At each fold, the training dataset is divided into two subsets, ‘search’, and ‘training’, by a 70/30 fraction. The ‘search’ set is used in the architecture search; the ‘training’

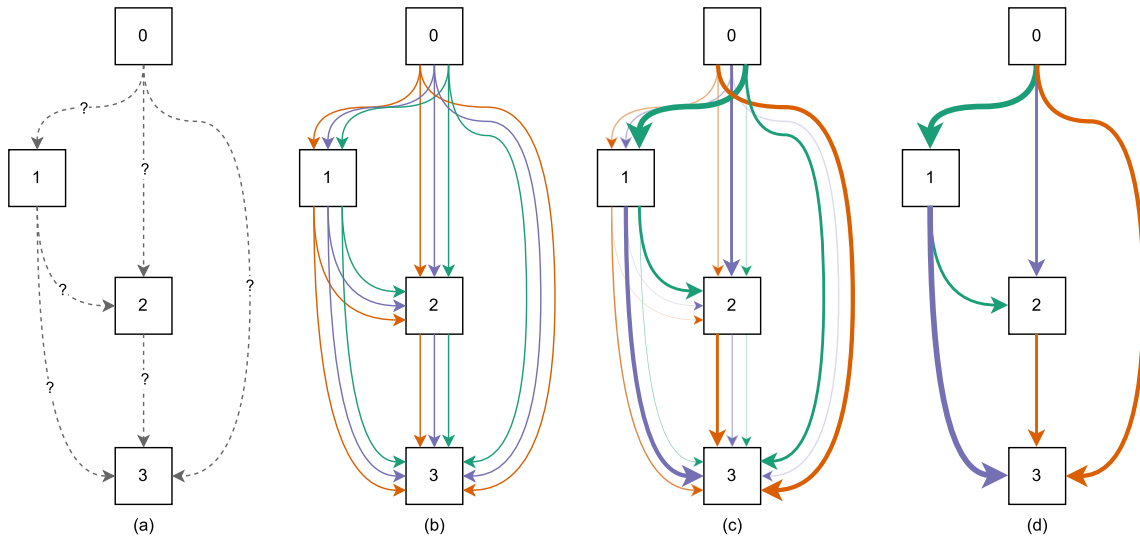


FIGURE 3. DARTS employs steps (a) to (d) to search cell architectures: (a) initialises the graph, (b) forms a search space, (c) updates edge weights, and (d) determines the final cell structure. Nodes signify representations, edges represent operations, with light-coloured edges indicating weaker and dark-coloured edges representing stronger operations.

set is used in optimising the searched architecture, and the remaining testing dataset is used to infer and obtain the testing performance of the searched and optimised model. This way, we manage to split the dataset into three sets in each cross-validation session. The IEMOCAP dataset has five sessions with ten actors and two unique speakers in each. We use one session for the testing dataset and four sessions for the search and training datasets. Similarly, MSP-Improv comprises six sessions including twelve actors. We take one session in the testing dataset and the remaining five sessions in the search and training dataset. MSP-Podcast includes a speaker ID with each audio utterance, and we group the entire dataset by the speaker and divide it by the 70/30 rule.

In this research, we use Mel Frequency Cepstral Coefficients (MFCC) as input features to the model. MFCC has been used as the input feature in many SER studies in the literature [34], [35] and has proven to obtain promising results. Some machine learning research uses the Mel Filter bank as an input feature when the algorithm is not vulnerable to strongly correlated data. We picked the MFCC for this study since the deep learning model is produced automatically and we do not want to infer the model’s sensitivity to correlated input. We extract 128 MFCCs from each 8-second audio utterance from the dataset. If the audio utterance length is less than 8 seconds, we added padding with zeros while the lengthier utterances are truncated. The MFCC extraction from the Librosa python library [36] outputs a shape 128×512 , downsampled with max pooling, to create a spectrogram of the shape 128×128 .

B. BASELINE MODELS

We compare the performance of emoDARTS for SER with three models developed without DARTS (w/o DARTS): 1) CNN, 2) CNN+ LSTM, and 3) CNN+LSTM with




FIGURE 4. Visualisation of the CNN+LSTM attention baseline model. The parameters of the CNN layer are: kernel size (k)=2, stride (s)=2 and, padding (p)=2 and the parameters of the Max-pooling layer are: kernel size (k)=2 and stride (s)=2 and the LSTM layer has 128 units.

attention as baseline models. The CNN baseline model consists of a CNN layer (kernel size=2, stride=2, and padding=2) followed by a Max-Pooling layer (kernel size=2 and stride=2). Two dense layers then processes the output from the Max-Pooling layer after applying a dropout of 0.3. Finally, the last dense layer has four output units resembling the four emotion classes, and the model outputs the probability estimation of each emotion for a given input by a Softmax function.

The CNN+LSTM baseline model is built, including an additional bi-directional LSTM layer of 128 units after the Max-Pooling layer. An attention layer is added to the LSTM layer in the ‘CNN+LSTM attention’ baseline model. Figure 4 shows the architecture of the CNN+LSTM attention baseline model.

C. DARTS CONFIGURATION

We divide the cell search space operations into the two separate parts CNN and SeqNN based on the components they apply. Table 2 lists the type of operations used in each component. The cell search space of the CNN component consists of pooling operations such as 3×3 max pooling

TABLE 2. Type of DARTS operations used in each component. “ i ” represents the kernel size in CNN while “ j ” represents the number of layers in the SeqNN component.

Component	Operation	Description
CNN	max_pool_ ixi	Max Pooling layer with kernel i
CNN	avg_pool_ ixi	Average Pooling layer with kernel i
CNN	dil_conv_ ixi	Dilated Convolution layer with kernel i , and dilation 2
CNN	sep_conv_ ixi	Two Convolution layers with kernel i
CNN	conv_ $ix1_1xi$	Two Convolution layers with first kernel ($ix1$) and second ($1xi$)
SeqNN	lstm_ j	LSTM with j layers
SeqNN	lstm_att_ j	LSTM of j layers with Attention
SeqNN	rnn_ j	RNN with j layers
SeqNN	rnn_att_ j	RNN of j layers with Attention

($i = 3$) and 3×3 average pooling ($i = 3$), convolutional operations such as 3×3 and 5×5 separable convolutions ($i = 3, 5$), 3×3 and 5×5 dilated convolution ($i = 3, 5$), $7 \times 1 - 1 \times 7$ factorised convolution ($i = 7$), identity connections, and no connections while the SeqNN component consists of operations such as RNN of layers 1 through 4 ($j = 1, 2, 3, 4$), RNN of layers 1 and 2 with attention ($j = 1, 2$), LSTM of layers 1 through 4 ($j = 1, 2, 3, 4$), LSTM of layers 1 and 2 with attention ($j = 1, 2$), identity connections and no connections.

We use stochastic gradient descent with a learning rate from 0.025 to 0.001 using a cosine annealing schedule as the optimiser to optimise the weights of the operations. The search is run for 300 epochs.

In our experiments, we use four DARTS cells ($C = 4$) for the CNN component following the work of Liu et al. [20] and two DARTS cells ($N = 2$) for the SeqNN component. The intuition of using $N = 2$ for the SeqNN component is discussed in section VI-B. As defined in [20], we apply reduction cells at every $\frac{1}{3}C^{th}$ and $\frac{2}{3}C^{th}$ position of the layers in CNN component. We randomly initialise α values and the DARTS search algorithm optimises α values related to each operation. The output from the CNN component is flattened to a vector before passing to the SeqNN component to adjust the input dimension of the RNN and LSTM layers.

Once the search operation completes, it outputs the architecture of DARTS cells, which is called “genome”. We create a deep learning model with the CNN component having four CNN cells and the SeqNN component having two SeqNN cells. This model is trained for 300 epochs with the training set of the datasets to minimise the loss.

We use the popular deep learning library PyTorch [37] for model development and training. The experiments are run on an NVIDIA A100 GPU with 40 GB of VRAM. We published the source code related to our research in a dedicated GitHub repository, allowing for smooth replication of our research findings.¹

¹<https://github.com/iot-health/emoDARTS>

V. EVALUATION

We report the results using the Unweighted Accuracy (UA%), calculated by dividing the total of all classes’ recall by their number. This is recognised to depict unbalanced data workloads intrinsic to SER accurately. We additionally provide the Weighted Accuracy (WA%) mainly to compare our results with relevant studies [9], [10]. Last but not least, we also report the number of parameters of the model as an indication of the model’s complexity, calculated by adding all trainable parameters in the created model.

A. CNN ONLY MODEL

We initially assess the performance of the CNN-only model generated by DARTS (CNN – DARTS) compared to our benchmark model, specifically CNN – w/o DARTS, using the IEMOCAP dataset. The results, detailed in Table 3, reveal that the DARTS-generated CNN model outperforms the performance of the baseline SER model. Additionally, Table 3 illustrates the performance of the DARTS-generated model with eight cells ($C = 8$), showing a lower performance compared to its counterpart with $C = 4$. This decline in performance with an increased number of cells indicates a rise in the model’s complexity, leading to overfitting and subsequent accuracy reduction.

We further examine the results from the CNN branch of Sun et al.’s EmotionNAS model [9] to highlight performance enhancements. For a direct and clear comparison of performance, we specifically utilise the ‘Spectrogram Branch’ of EmotionNAS, contrasting it with our ‘CNN – DARTS’ model. This focused comparison is chosen to ensure a fair evaluation since both models share a similar architecture. We see that the performance of our CNN – DARTS models surpasses the performance of the CNN branch of EmotionNAS by at least 5%. It is worth noting that the ‘whole model’ of EmotionNAS has a different architecture, employing a branched structure, while emoDARTS utilises a stacked architecture.

TABLE 3. Performance comparison between the DARTS generated CNN model (CNN – DARTS) and a CNN SER model developed without DARTS (CNN – w/o DARTS) for the IEMOCAP dataset. The number of parameters is in thousands.

Model	Param.	Cell	UA (%)	WA (%)
CNN – DARTS	417K	4	69.36 ± 3.00	72.55 ± 3.70
CNN – DARTS	428K	8	62.25 ± 6.74	63.78 ± 6.83
CNN – w/o DARTS	35K	-	50.04 ± 2.69	51.01 ± 2.23
EmotionNAS [CNN] [9]	130K	3	57.3	63.2

B. emoDARTS MODEL

We analyse the performance of the CNN-SeqNN model generated by DARTS (emoDARTS) in contrast to the SER models optimised without DARTS (w/o DARTS) and visualise this in Figure 5 and Table 4. The graph shows that the NAS-generated SER model performs better than the

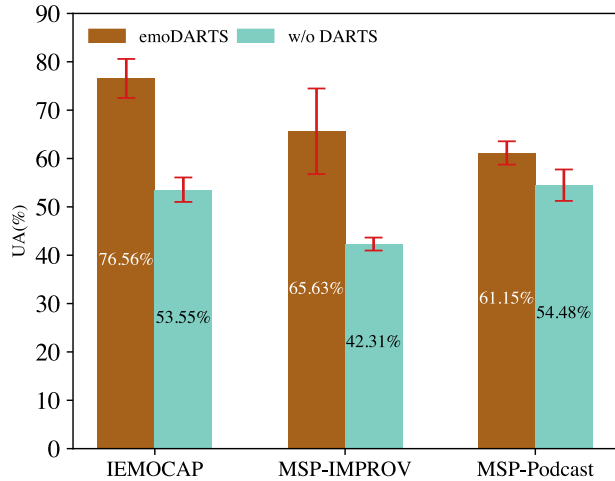


FIGURE 5. Comparison of UA% between the datasets the NAS generated (emoDARTS) and CNN+LSTM attention models developed without DARTS (w/o DARTS).

TABLE 4. Unweighted Accuracy (UA%) of the CNN+LSTM attention model developed without DARTS (w/o DARTS), Unweighted Accuracy (UA%), and Weighted Accuracy (WA%) of the emoDARTS model for each Dataset.

Dataset	w/o DARTS	emoDARTS	
	UA%	UA%	WA%
IEMOCAP	53.55 ± 2.53	76.56 ± 4.03	78.03 ± 3.51
MSP-IMPROV	42.31 ± 1.34	65.63 ± 8.85	65.32 ± 8.73
MSP-Podcast	54.48 ± 3.25	61.15 ± 2.41	62.33 ± 1.80

TABLE 5. Accuracy of SER models published by related studies compared with our study for the improvised subset in the IEMOCAP dataset.

Study	UA %	WA %
Sun H. et al. 2022 [9]	69.1	72.1
Wu X. et al. 2022 [10]	56.28	68.87
emoDARTS (Our Study)	76.56 ± 4.03	78.03 ± 3.51

baseline SER model developed without DARTS for the three datasets.

We also compare the performance of the SER model generated by our approach with the most related studies, ‘EmotionNAS’ of Sun H. et al., and the ‘CNN_RNN_att’ system of Wu X. et al. in Table 5. It is visible that the SER model generated by our methodology for the IEMOCAP dataset outperforms the SER models generated by DARTS in the related literature.

It is further worthwhile to investigate the rationale for increased performance when compared to the results of Wu et al.’s ‘CNN RNN att’ [10] system. We suggest the improved performance is due to the relaxed candidate operations order rather than the pre-defined layer order. In Wu et al.’s study [10], for example, the initial layers are pre-defined to be convolutional layers. The DARTS algorithm must select the best convolutional layer from a pool of just CNN layers. In contrast, our technique allows DARTS to choose among many operations such as convolutions,

pooling, and skip connections. Figure 6 shows one such use scenario, in which the DARTS searched architecture consists of pooling layers in the initial segments.

Figure 6 shows a visualisation of the architecture for each type of cell (normal (a) and reduction (b) cell of the CNN component, and cell in the SeqNN component (c) searched by DARTS for the emoDARTS model. It is visible that DARTS has selected three LSTM based operations for the SeqNN component and only one of them contains attention. This shows that jointly optimising the emoDARTS model has enabled the DARTS framework to choose optimum operations rather than blindly choosing layers with ‘attention’ for all the operations.

C. RESTRICTING THE SEARCH SCOPE

We study the impact on the performance of the searched model by restricting the search scope for the SeqNN component. We divide the search scope into five segments namely ‘LSTM Only’, ‘LSTM-Att. Only’, ‘RNN Only’, and ‘RNN-Att. Only’. Table 6 shows the DARTS operations allowed as the candidate operations in the SeqNN component during the search phase.

Table 7 shows the performance and number of parameters of the searched model when the candidate search operations are restricted. Here, we study the effect on the performance of the searched architecture when the search algorithm was only given a restricted set of operations. For example, the ‘LSTM Only’ study only allowed to use operations from lstm_1, lstm_2, lstm_3, and lstm_4. We try to identify the most important types of genome operations that we can use in the search algorithm. This approach allows to use of only the important operations in the search scope and optimises the memory utilisation in the search phase.

TABLE 6. DARTS operations allowed as the candidate operations in the SeqNN component during the search phase.

Scope	Candidate Operations
emoDARTS	lstm_1, lstm_2, lstm_3, lstm_4, lstm_att_1, lstm_att_2, rnn_1, rnn_2, rnn_3, rnn_4, rnn_att_1, rnn_att_2
LSTM Only	lstm_1, lstm_2, lstm_3, lstm_4
LSTM-Att. Only	lstm_att_1, lstm_att_2
RNN Only	rnn_1, rnn_2, rnn_3, rnn_4
RNN-Att. Only	rnn_att_1, rnn_att_2

Comparing the trials ‘emoDARTS’ and ‘LSTM Only’ in the IEMOCAP dataset, we can observe that even though the number of parameters has tripled in the ‘LSTM Only’ scenario, the performance (UA%) has not increased. This indicates that increasing the number of parameters just by increasing the complexity of the model does not tend to give better performance, but the model components should be compatible with each other.

Notably, models using ‘RNN Only’ genomic operations achieve the second-highest accuracy despite having

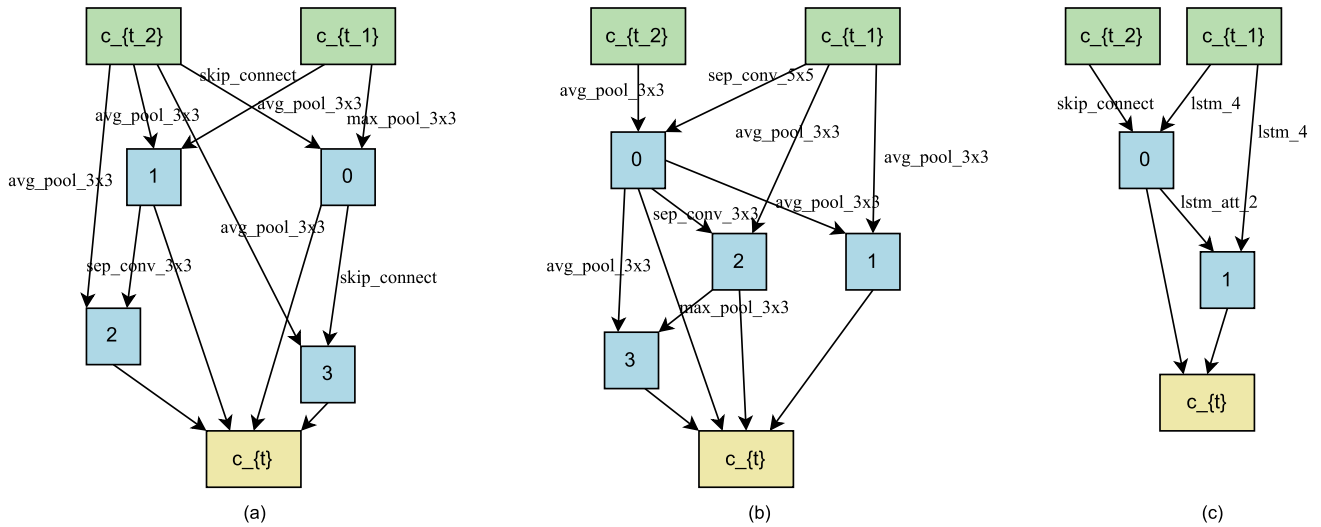


FIGURE 6. DARTS searched t^{th} cell structure for the CNN Normal Cell (a), CNN Reduction Cell (b), and SeqNN cell (c) for the emoDARTS model.

TABLE 7. Performance (UA%, WA%) and number of parameters (Param.) of each generated model when the candidate search operations are restricted.

Genome Ops.	IEMOCAP			MSP-Improv			MSP-Podcast		
	UA%	WA%	Param.	UA%	WA%	Param.	UA%	WA%	Param.
emoDARTS	76.55 ± 4.03	78.03 ± 3.51	1 014 556	65.63 ± 8.85	65.32 ± 8.73	1 008 812	61.15 ± 2.41	62.33 ± 1.79	1 835 604
LSTM Only	67.42 ± 5.76	67.78 ± 5.97	2 939 372	32.35 ± 1.93	31.85 ± 0.98	2 931 748	53.12 ± 5.82	54.48 ± 5.43	2 943 852
LSTM-Att. Only	68.12 ± 7.12	68.59 ± 6.62	7 284 412	49.74 ± 10.08	49.8 ± 10.15	6 474 676	53.11 ± 5.83	55.07 ± 5.14	7 284 932
RNN Only	70.42 ± 8.32	72.07 ± 8.23	356 492	63.59 ± 7.88	63.49 ± 7.95	1 016 284	59.16 ± 6.67	60.72 ± 7.65	706 068
RNN-Att. Only	36.14 ± 10.97	48.63 ± 6.56	29 324	57.71 ± 12.61	57.04 ± 12.14	22 932	59.29 ± 6.38	60.25 ± 6.09	2 844 420

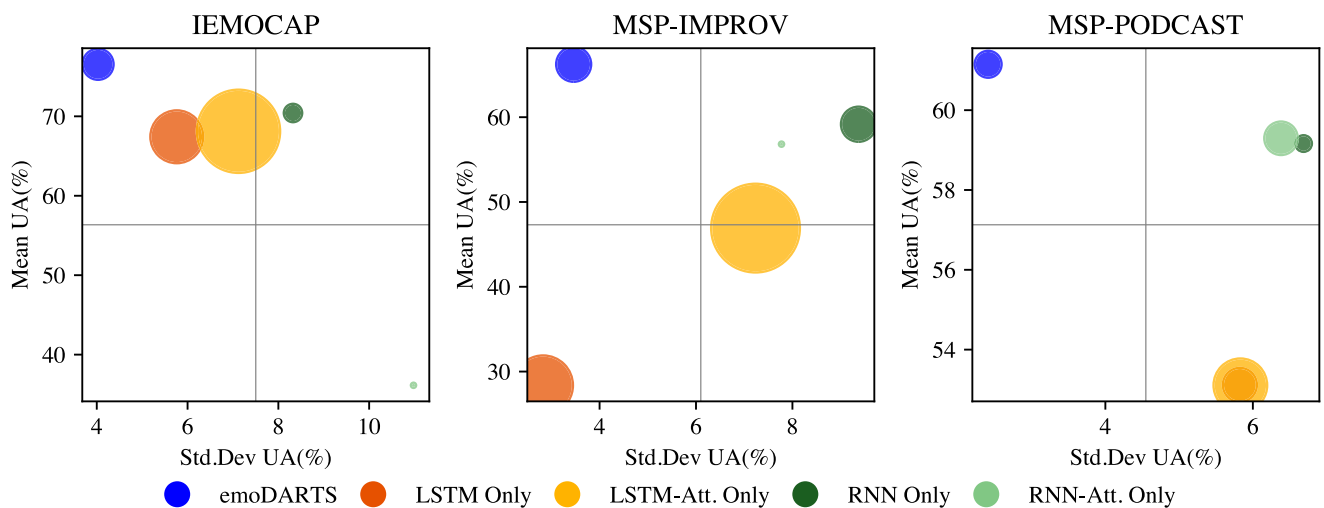


FIGURE 7. Visualisation of results for the studies restricting the search space for the three datasets: IEMOCAP, MSP-IMPROV, and MSP-Podcast. The vertical axis is the mean UA% and the horizontal axis is the standard deviation of UA%. The size of the marker depicts the size (number of parameters) of the generated model. The best performing model can be found at the top-left most position of the figure which has the highest mean UA% and lowest standard deviation of UA%.

much fewer trainable parameters. Figure 8 depicts the cell architecture, which consists mostly of pooling layers and skip connection operations that do not have any training parameters and hence do not contribute to the total number of trainable parameters.

We provide in Table 7 results as well as a scatter plot for better visualisation in Figure 7, where the mean UA% in the vertical axis, standard deviation of UA% in the horizontal axis and size of the markers indicates the number of parameters. The plots indicate which model gives better performance in

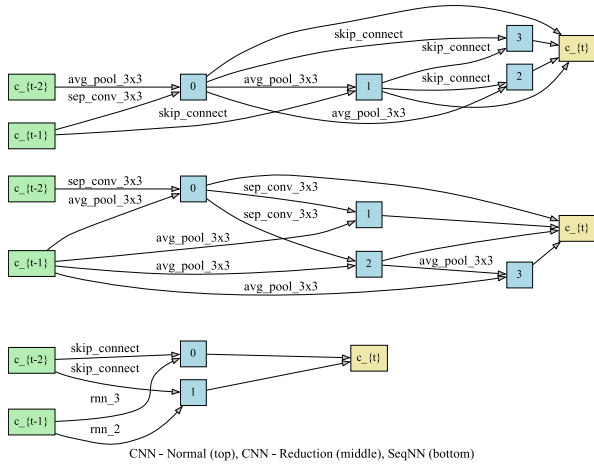


FIGURE 8. DARTS searched cell structure for CNN and SeqNN cells when the SeqNN search space has only RNN operations (RNN Only) for the IEMOCAP dataset.

terms of the mean accuracy and its standard deviation. The best performing model can be found in the top left corner of the plot where the highest UA% and lowest Standard deviation of UA% are present. According to the figure, the best performing model for all three datasets is given by ‘emoDARTS’ where all the candidate operations are available in the search scope.

VI. DISCUSSION

Throughout this study, we encountered various challenges. In this section, we report the key challenges and our strategies for overcoming them. The three primary challenges we faced were:

- 1) Optimising the GPU Memory utilisation
- 2) Converging to a local minima
- 3) High Standard Deviation of the results

A. OPTIMIZING THE GPU MEMORY UTILIZATION

The DARTS algorithm conceptualises the search problem as a network graph, establishing multiple edges between each node. The quantity of edges corresponds to the defined candidate operations. These operations encompass various possibilities, ranging from simple CNN, pooling, and RNN layers to intricate modules like an LSTM-attention module. In the search phase, a super-neural network is constructed, resulting in multiple instances of neural network layers or modules within this overarching structure.

Figure 9 shows an example of a graph inside a DARTS cell that has four nodes and the candidate operations are “lstm_1”, “lstm_2”, and “lstm_att_1”, where “lstm_1” is a single layer LSTM component, “lstm_2” is a double layer LSTM component, and “lstm_att_1” is an attention induced single layer LSTM component. According to the example, a single DARTS cell should initiate 6×lstm_1 layers, 6×lstm_2 layers, and 6×lstm_att_1 layers. If the search configuration has 4 cells, we have to initialise




FIGURE 9. Example graph of a DARTS cell which has four nodes and candidate operations are “lstm_1”, “lstm_2”, and “lstm_att_1”. The same edge colour denotes the same type of operation.

4 instances of cells where all the weight and bias parameters have to be initialised in the computing device. This will increase the GPU memory utilisation.

Providing a higher number of nodes in a cell, a higher number of cells, and expanding the array of candidate operations will increase the amount of GPU memory utilisation and eventually will exhaust the GPU memory capacity failing the search operation.

To optimise the GPU memory utilisation, we recommend conducting an assessment to determine the set of possible search operations and hyperparameters such as the number of layers, cells, and nodes inside the cell considering the GPU resources available.

On the other hand, based on the results of Table 7, we should not be restricted only to a single type of network but rather should consist of a variety of network architectures. We selected the set of candidate operations indicated in Table 6 under ‘emoDARTS’ based on GPU resource availability and on the premise that all sorts of candidate operations should be available in the search space.

B. CONVERGING TO A LOCAL MINIMA

Throughout the course of our experiments, we attempted various configurations for the number of cells and nodes. We observed that the SeqNN module converges to a local minimum when the number of cells and number of nodes is greater than 3. The output of the searched genome for the SeqNN module contained all “skip_connect” which indicates identity operations are used instead of any RNN or LSTM operations. Figure 10 shows one such instance DARTS SeqNN genome.

We were able to address the challenge by reducing the complexity of the candidate search graph by reducing the number of cells and the number of nodes inside a cell. More research is, however, needed to manage a more complex search network.

C. HIGH STANDARD DEVIATION IN THE RESULTS

An important observation derived from our results is the high standard deviation. This can be attributed to the dataset-splitting method we employed. Specifically, we adopt speaker-independent dataset splitting, where the training and validation sets are segregated based on the speaker. In this configuration, any audio utterance from a particular speaker




FIGURE 10. DARTS genome of the SeqNN module where the search algorithm selected the identity operation as all the operations.

in the validation set remains unseen by the model during training. Consequently, the DARTS-optimised model is not trained to handle the data distribution of the validation set. To tackle this challenge, potential solutions include dataset poisoning and enhancing the generalisation capabilities of the SER model by incorporating dropout layers.

VII. CONCLUSION

In conclusion, this paper introduced an innovative approach to enhancing speech emotion recognition (SER) using differentiable architecture search (DARTS). Our primary focus was on tailoring DARTS for a joint configuration of a Convolutional Neural Network (CNN) and a Sequential Neural Network, deviating from previous studies by allowing DARTS to autonomously determine the optimal layer order for the CNN within the DARTS cell without imposing constraints.

A comprehensive evaluation was conducted, comparing our proposed method with baseline models developed without DARTS and various genome operations, including LSTM only, LSTM with attention only, RNN only, and RNN with attention only. The detailed assessments consistently demonstrate the superior performance of our proposed method. Contrasting with existing studies further validates the effectiveness and superiority of our approach, considering parameter size and accuracy as essential dimensions for comparison.

Notably, our study extends beyond the confines of the commonly used IEMOCAP dataset, incorporating two additional datasets, MSP-IMPROV and MSP-Podcast. This extension showcases the superior performance of our proposed method across diverse datasets, affirming its generalisation capability.

Furthermore, we shared valuable insights gained from our experiences, addressing challenges related to GPU exhaustion and converging to local minima. These insights serve as practical guidance for researchers, helping them navigate potential pitfalls and optimise the application of DARTS in SER.

Future efforts will need to deal with neural architecture for further modern architectures such as transformers and translating the made findings beyond the targeted field of application.

REFERENCES

- [1] J. Zhao, X. Mao, and L. Chen, "Speech emotion recognition using deep 1D & 2D CNN LSTM networks," *Biomed. Signal Process. Control*, vol. 47, pp. 312–323, Jan. 2019.
- [2] M. A. Jalal, R. Milner, and T. Hain, "Empirical interpretation of speech emotion perception with attention based model for speech emotion recognition," in *Proc. Annu. Conf. Int. Speech Commun. Assoc. (INTERSPEECH)*, Oct. 2020, pp. 4113–4117.
- [3] E. Lieskovská, M. Jakubec, R. Jarina, and M. Chmulk, "A review on speech emotion recognition using deep learning and attention mechanism," *Electronics*, vol. 10, no. 10, p. 1163, May 2021.
- [4] S. Latif, R. Rana, S. Khalifa, R. Jurdak, and B. W. Schuller, "Multitask learning from augmented auxiliary data for improving speech emotion recognition," *IEEE Trans. Affect. Comput.*, vol. 14, no. 4, pp. 3164–3176, Oct. 2023. [Online]. Available: <https://ieeexplore.ieee.org/document/9947296/>
- [5] W. Han, H. Ruan, X. Chen, Z. Wang, H. Li, and B. Schuller, "Towards temporal modelling of categorical speech emotion recognition," in *Proc. Annu. Conf. Int. Speech Commun. Assoc. (INTERSPEECH)*, Sep. 2018, pp. 932–936.
- [6] K. N. Haque, M. A. Yousuf, and R. Rana, "Image denoising and restoration with CNN-LSTM encoder decoder with direct attention," 2018, *arXiv:1801.05141*.
- [7] Y. Li, T. Zhao, and T. Kawahara, "Improved end-to-end speech emotion recognition using self attention mechanism and multitask learning," in *Proc. Annu. Conf. Int. Speech Commun. Assoc. (INTERSPEECH)*, Sep. 2019, pp. 2803–2807.
- [8] S. Latif, R. Rana, S. Khalifa, R. Jurdak, and B. W. Schuller, "Self supervised adversarial domain adaptation for cross-corpus and cross-language speech emotion recognition," *IEEE Trans. Affect. Comput.*, vol. 14, no. 3, pp. 1912–1926, Jul./Sep. 2023.
- [9] H. Sun, Z. Lian, B. Liu, Y. Li, J. Tao, L. Sun, C. Cai, M. Wang, and Y. Cheng, "EmotionNAS: Two-stream neural architecture search for speech emotion recognition," in *Proc. Annu. Conf. Int. Speech Commun. Assoc. (INTERSPEECH)*, Aug. 2023, pp. 3597–3601.
- [10] X. Wu, S. Hu, Z. Wu, X. Liu, and H. Meng, "Neural architecture search for speech emotion recognition," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2022, pp. 6902–6906.
- [11] Y. Xu, L. Xie, X. Zhang, X. Chen, G.-J. Qi, Q. Tian, and H. Xiong, "PC-DARTS: Partial channel connections for memory-efficient architecture search," in *Proc. Int. Conf. Learn. Represent.*, 2020.
- [12] G. M. Biju, G. N. Pillai, and J. Seshadrinath, "Electric load demand forecasting with RNN cell generated by DARTS," in *Proc. IEEE Region 10 Conf. (TENCON)*, Oct. 2019, pp. 2111–2116.
- [13] R. Maulik, R. Egele, B. Lusch, and P. Balaprakash, "Recurrent neural network architecture search for geophysical emulation," in *Proc. Int. Conf. High Perform. Comput., Netw., Storage Anal.*, Nov. 2020, pp. 1–14.
- [14] W. Q. Zheng, J. S. Yu, and Y. X. Zou, "An experimental study of speech emotion recognition based on deep convolutional neural networks," in *Proc. Int. Conf. Affect. Comput. Intell. Interact. (ACII)*, Sep. 2015, pp. 827–831.
- [15] G. Trigeorgis, F. Ringeval, R. Brueckner, E. Marchi, M. A. Nicolaou, B. Schuller, and S. Zafeiriou, "Adieu features? End-to-end speech emotion recognition using a deep convolutional recurrent network," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Mar. 2016, pp. 5200–5204.
- [16] Z. Zhang, J. Han, K. Qian, and B. Schuller, "Evolving learning for analysing mood-related infant vocalisation," in *Proc. Annu. Conf. Int. Speech Commun. Assoc. (INTERSPEECH)*, Sep. 2018, pp. 142–146.
- [17] B. Zoph and Q. V. Le, "Neural architecture search with reinforcement learning," in *Proc. 5th Int. Conf. Learn. Represent. (ICLR)*, 2017.
- [18] E. Liberis, L. Dudziak, and N. D. Lane, " μ NAS: Constrained neural architecture search for microcontrollers," in *Proc. 1st Workshop Mach. Learn. Syst.*, Apr. 2021, pp. 70–79.

- [19] C. Gong, Z. Jiang, D. Wang, Y. Lin, Q. Liu, and D. Z. Pan, "Mixed precision neural architecture search for energy efficient deep learning," in *Proc. IEEE/ACM Int. Conf. Computer-Aided Design (ICCAD)*, Nov. 2019, pp. 1–7.
- [20] H. Liu, K. Simonyan, and Y. Yang, "DARTS: Differentiable architecture search," 2018, *arXiv:1806.09055*.
- [21] C. Liu, B. Zoph, M. Neumann, J. Shlens, W. Hua, L.-J. Li, L. Fei-Fei, A. Yuille, J. Huang, and K. Murphy, "Progressive neural architecture search," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 19–34.
- [22] H. Liang, S. Zhang, J. Sun, X. He, W. Huang, K. Zhuang, and Z. Li, "DARTS+: Improved differentiable architecture search with early stopping," 2019, *arXiv:1909.06035*.
- [23] A. Wan, X. Dai, P. Zhang, Z. He, Y. Tian, S. Xie, B. Wu, M. Yu, T. Xu, K. Chen, P. Vajda, and J. E. Gonzalez, "FBNetV2: Differentiable neural architecture search for spatial and channel dimensions," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 12962–12971.
- [24] X. Chen and C.-J. Hsieh, "Stabilizing differentiable architecture search via perturbation-based regularization," in *Proc. 37th Int. Conf. Mach. Learn.*, 2020, pp. 1554–1565.
- [25] H. Yu, H. Peng, Y. Huang, J. Fu, H. Du, L. Wang, and H. Ling, "Cyclic differentiable architecture search," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 1, pp. 211–228, Jan. 2023.
- [26] T. Rajapakshe, R. Rana, S. Khalifa, B. Sisman, and B. Schuller, "Enhancing speech emotion recognition through differentiable architecture search," 2023, *arXiv:2305.14402*.
- [27] Y. Zhang, J. Du, Z. Wang, J. Zhang, and Y. Tu, "Attention based fully convolutional network for speech emotion recognition," in *Proc. Asia-Pacific Signal Inf. Process. Assoc. Annu. Summit Conf. (APSIPA ASC)*, Nov. 2018, pp. 1771–1775.
- [28] M. Chen, X. He, J. Yang, and H. Zhang, "3-D convolutional recurrent neural networks with attention model for speech emotion recognition," *IEEE Signal Process. Lett.*, vol. 25, no. 10, pp. 1440–1444, Oct. 2018.
- [29] H. Zou, Y. Si, C. Chen, D. Rajan, and E. S. Chng, "Speech emotion recognition with co-attention based multi-level acoustic information," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2022, pp. 7367–7371.
- [30] Z.-T. Liu, M.-T. Han, B.-H. Wu, and A. Rehman, "Speech emotion recognition based on convolutional neural network with attention-based bidirectional long short-term memory network and multi-task learning," *Appl. Acoust.*, vol. 202, Jan. 2023, Art. no. 109178.
- [31] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "IEMOCAP: Interactive emotional dyadic motion capture database," *Lang. Resour. Eval.*, vol. 42, no. 4, pp. 335–359, Dec. 2008.
- [32] C. Busso, S. Parthasarathy, A. Burmania, M. AbdelWahab, N. Sadoughi, and E. M. Provost, "MSP-IMPROV: An acted corpus of dyadic interactions to study emotion perception," *IEEE Trans. Affect. Comput.*, vol. 8, no. 1, pp. 67–80, Jan. 2017.
- [33] R. Lotfian and C. Busso, "Building naturalistic emotionally balanced speech corpus by retrieving emotional speech from existing podcast recordings," *IEEE Trans. Affect. Comput.*, vol. 10, no. 4, pp. 471–483, Oct. 2019.
- [34] S. Davis and P. Mermelstein, "Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-28, no. 4, pp. 357–366, Aug. 1980.
- [35] S. Latif, R. Rana, S. Khalifa, R. Jurdak, and J. Epps, "Direct modelling of speech emotion from raw speech," in *Proc. Annu. Conf. Int. Speech Commun. Assoc. (INTERSPEECH)*, Sep. 2019, pp. 3920–3924.
- [36] B. McFee, C. Raffel, D. Liang, D. Ellis, M. McVicar, E. Battenberg, and O. Nieto, "Librosa: Audio and music signal analysis in Python," in *Proc. 14th Python Sci. Conf.*, vol. 8, 2015, pp. 18–24.
- [37] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in *Proc. Adv. Neural Inf. Process. Syst.*, Dec. 2019, pp. 8024–8035.




THEJAN RAJAPAKSHE (Student Member, IEEE) received the bachelor's degree in applied sciences from the Rajarata University of Sri Lanka and the bachelor's degree in information technology from the School of Computing, University of Colombo, in 2016.

He is currently a Research Scholar with the University of Southern Queensland (UniSQ) and a Consultant at Cognizant Servian. His research interests include reinforcement learning, speech processing, and deep learning.



RAJIB RANA (Member, IEEE) received the B.Sc. degree in computer science and engineering from Khulna University, Bangladesh, and the Ph.D. degree in computer science and engineering from the University of New South Wales, Sydney, Australia, in 2011.

He is currently an Experimental Computer Scientist, an Advance Queensland Senior Research Fellow, and an Associate Professor of computer science with the University of Southern Queensland. He is also the Director of the IoT Health Research Program, University of Southern Queensland. He received his postdoctoral training at the Autonomous Systems Laboratory, CSIRO, before joining the University of Southern Queensland, in 2015. His research work aims to capitalize on advancements in technology along with sophisticated information and data processing to better understand disease progression in chronic health conditions and develop predictive algorithms for chronic diseases, such as mental illness and cancer. His current research interests include unsupervised representation learning, reinforcement learning, adversarial machine learning, and emotional speech generation. He was a recipient of the prestigious Young Tall Poppy QLD Award, in 2018, as one of Queensland's most outstanding scientists for achievements in the area of scientific research and communication. He was also a recipient of the Prime Minister and President's Gold Medal for outstanding achievements.



SARA KHALIFA (Member, IEEE) received the Ph.D. degree in computer science and engineering from the University of New South Wales (UNSW), Sydney, NSW, Australia.

She is currently an Associate Professor with the Queensland University of Technology. She is also an Honorary Adjunct Lecturer with The University of Queensland and a Conjoint Lecturer with UNSW. Her research interests include the broad aspects of mobile and ubiquitous computing, mobile sensing, and Internet of Things (IoT). Her Ph.D. dissertation received the 2017 John Makepeace Bennett Award which is awarded by the Computing Research and Education Association of Australasia (CORE) to the Best Ph.D. Dissertation of the year within Australia and New Zealand in the field of computer science. Her research has been recognized by multiple awards, including the 2017 NSW Mobility Innovation of the year, the 2017 NSW Research and Development Innovation of the year, National Merit Research and Development Innovation of the year, and the Merit Research and Development Award at the Asia Pacific ICT Alliance (APICTA) Awards, commonly known as the 'Oscar' of the ICT industry in the Asia-Pacific, among others.



BERRAK SISMAN (Member, IEEE) received the Ph.D. degree in electrical and computer engineering from the National University of Singapore, in 2020, fully funded by A*STAR Graduate Academy under Singapore International Graduate Award (SINGA). She was a Visiting Researcher with the RIKEN Advanced Intelligence Project, Japan, in 2018. She was an Exchange Ph.D. Student with the University of Edinburgh and a Visiting Scholar with The Centre for Speech Technology Research (CSTR), The University of Edinburgh, in 2019. She was a Postdoctoral Research Fellow with the National University of Singapore, from 2019 to 2020. She is currently a tenure-track Assistant Professor with the Erik Jonsson School of Electrical and Computer Engineering, The University of Texas at Dallas (UT Dallas), USA. Prior to joining UT Dallas, she was a Faculty Member with Singapore University of Technology and Design, from 2020 to 2022, where she taught machine learning and deep learning courses. Her research interests include machine learning, signal processing, emotion, speech synthesis, and voice conversion.



BJÖRN W. SCHULLER (Fellow, IEEE) received the Diploma degree in electrical engineering and information technology, the Ph.D. degree in automatic speech and emotion recognition (electrical engineering and information technology), and the Habilitation and Adjunct Teaching Professorship degrees in signal processing and machine intelligence (electrical engineering and information technology) from TUM, Munich, Germany, in 1999, 2006, and 2012, respectively.

He is currently a Professor of artificial intelligence with the Department of Computing, Imperial College London, U.K., where he heads the Group on Language, Audio and Music (GLAM), a Full Professor and the Head of the Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, Germany, and the CEO/CSO of audeERING. Previously, he was a Full Professor and the Head of the Chair of Complex and Intelligent Systems, University of Passau, Germany. He (co-)authored five books and more than 1200 publications in peer-reviewed books, journals, and conference proceedings leading to more than overall 48000 citations (H-index of 99). He is a Golden Core Member of the IEEE Computer Society, a fellow of the ISCA and BCS, a Senior Member of the ACM, and a Fellow and President-emeritus of the Association for the Advancement of Affective Computing (AAAC). He was the General Chair of ACII 2019, the Program Chair of Interspeech 2019 and ICMI 2019, a repeated Area Chair of ICASSP, and the former Editor-in-Chief of IEEE TRANSACTIONS ON AFFECTIVE COMPUTING. He is the Field Chief Editor of the *Frontiers in Digital Health*.



CARLOS BUSSO (Fellow, IEEE) is currently a Professor with the Department of Electrical and Computer Engineering, The University of Texas at Dallas, where he is also the Director of the Multimodal Signal Processing (MSP) Laboratory. He has worked on audio-visual emotion recognition, analysis of emotional modulation in gestures and speech, designing realistic human-like virtual characters, and detection of driver distractions. His research interests include human-centered multimodal machine intelligence and application, with a focus on the broad areas of affective computing, multimodal human-machine interfaces, in-vehicle active safety systems, and machine learning methods for multimodal processing. He is a member of the ISCA and AAAC, and a Senior Member of ACM. He was a recipient of the NSF CAREER Award. In 2014, he was a recipient of the ICMI Ten-Year Technical Impact Award. In 2015, his student received the Third Prize of the IEEE ITSS Best Dissertation Award (N. Li). He also received the Hewlett Packard Best Paper Award at the IEEE ICME 2011 (with J. Jain) and the Best Paper Award at the AAAC ACII 2017 (with Yannakakis and Cowie). He received the Best of IEEE TRANSACTIONS ON AFFECTIVE COMPUTING Paper Collection, in 2021 (with R. Lotfian) and the Best Paper Award from IEEE TRANSACTIONS ON AFFECTIVE COMPUTING, in 2022 (with Yannakakis and Cowie). He received the ACM ICMI Community Service Award, in 2023. In 2023, he received the Distinguished Alumni Award in the Mid-Career/Academia category by the Signal and Image Processing Institute (SIPI), University of Southern California. He is currently serving as an Associate Editor of IEEE TRANSACTIONS ON AFFECTIVE COMPUTING.

...