arXiv:2006.13211v1 [cs.CV] 23 Jun 2020

Meta Transfer Learning for Emotion Recognition

Dung Nguyen[a], Sridha Sridharan[a], Duc Thanh Nguyen[b], Simon Denman[a], David Dean[a], Clinton Fookes[a]

[a]*Speech, Audio, Image and Video Technology (SAIVT) Laboratory - Queensland University of Technology*
[b]*School of Information Technology - Deakin University*

Abstract

Deep learning has been widely adopted in automatic emotion recognition and has lead to significant progress in the field. However, due to insufficient annotated emotion datasets, pre-trained models are limited in their generalization capability and thus lead to poor performance on novel test sets. To mitigate this challenge, transfer learning performing fine-tuning on pre-trained models has been applied. However, the fine-tuned knowledge may overwrite and/or discard important knowledge learned from pre-trained models. In this paper, we address this issue by proposing a PathNet-based transfer learning method that is able to transfer emotional knowledge learned from one visual/audio emotion domain to another visual/audio emotion domain, and transfer the emotional knowledge learned from multiple audio emotion domains into one another to improve overall emotion recognition accuracy. To show the robustness of our proposed system, various sets of experiments for facial expression recognition and speech emotion recognition task on three emotion datasets: SAVEE, EMODB, and eNTERFACE have been carried out. The experimental results indicate that our proposed system is capable of improving the performance of emotion recognition, making its performance substantially superior to the recent proposed fine-tuning/pre-trained models based transfer learning methods.

Keywords: emotional knowledge transfer; emotion recognition; facial expression recognition; speech emotion recognition; transfer learning; cross-domain transfer; joint leaning

1. Introduction

Emotions of humans manifest in their facial expressions, voice, gestures, and posture. An accurate emotion recognition system based on one or a combination of these modalities would be useful in various applications including surveillance, medical, robotics, human computer interaction, affective computing, and automobile safety (Nguyen et al., 2017). Researchers in this area have focused mainly in the area of facial expression recognition to build reliable emotion recognition systems. This is still a challenging problem since very subtle emotional changes manifested in the facial expression could go undetected (Nguyen et al., 2017). Recently several approaches based on deep learning techniques have contributed to progressing in this area (Fan et al., 2016; Abbasnejad et al., 2017; Hasani and Mahoor, 2017).

In addition to facial expression stream, speech signals, which are regarded as one of the most natural media of human communication, carry both the contents of explicit linguistic and the information of implicit paralinguistic expressed by a speaker (Zhang et al., 2018). Due to this rich information contained in the speech, over the last two decades numerous studies and efforts have been devoted to progressing approaches, focusing on automatic and accurate detection of human emotions from speech signals (Zhang et al., 2018). Speech emotion recognition is presently playing an essential role in a wide range of applications such as automobile safety, surveillance, human computer interaction, and robotics, and is attracting a great deal of attention within the affective computing research community (Nguyen et al., 2018).

To develop solutions for speech emotion recognition, a number of methodologies have been proposed, in which researchers have primarily applied the use of hand engineered features based on the acoustic and paralinguistic information (Shen et al., 2011). Nevertheless, such hand-designed features seem not to be discriminative enough to boost the performance of speech emotion recognition (Zhang et al., 2018). Recently, algorithms based on deep learning techniques, which are capable of automatically learning features and also capable of modelling high-level information, have been the focus of most recent research and are gaining prominence.

Although the aforementioned deep learning approaches have made a great contri-

bution to progressing the emotion recognition area, we pointed out a key issue that plagues the advancement of emotion recognition research; e.g., the lack of sufficient quantities of annotated emotion data. This issue has become more critical with the advent of deep learning techniques which promise major improvements in emotion recognition accuracy in both single and multi-modal settings; yet we are unable to exploit the full potential of deep learning for emotion recognition due to the scarcity of annotated emotion data; deep learning techniques require large amounts of data for training. Transfer learning method, which commonly fine-tunes pre-trained/off-the-shelf CNN models on emotion dataset, have been widely investigated to overcome this problem. However, the representational features that are unrelated to emotion are still retained in off-the-shelf/pre-trained models and the extracted features are also vulnerable to identity variations in these approaches, leading to degrading the performance of the emotion recognition system fine-tuning off-the-shelf/pre-trained models on the emotion dataset.

To resolve these drawbacks, (Gideon et al., 2017) have exploited a progressive network originally proposed by (Rusu et al., 2016) which was able to potentially support transferring knowledge across sequences of tasks. (Gideon et al., 2017) have successfully transferred learning between three paralinguistic tasks: emotion, speaker, and gender recognition with an emphasis on speech emotion detection as the target application without the catastrophic forgetting effect. Their system outperformed the recent speech emotion recognition approaches utilizing fine-tuning pre-trained models and also performed significantly better than deep leaning models without the use of transfer learning techniques (Rusu et al., 2016). Nonetheless, an unavoidable limitation of this approach is that it is computationally intensive since a number of new networks keep on growing according to the demand for a increasing number of new tasks which need to be learned (Lee et al., 2017). More recently, in order to alleviate the aforementioned downsides, (Fernando et al., 2017) have proposed PathNet as an alternative novel learning algorithm for transfer learning. PathNet was designed as a neural network in which agents (e.g., pathways through different layers of the neural network) were embedded to discover which parts of the network to be re-used for new tasks (Fernando et al., 2017). Agents also hold an accountability for determining which subset

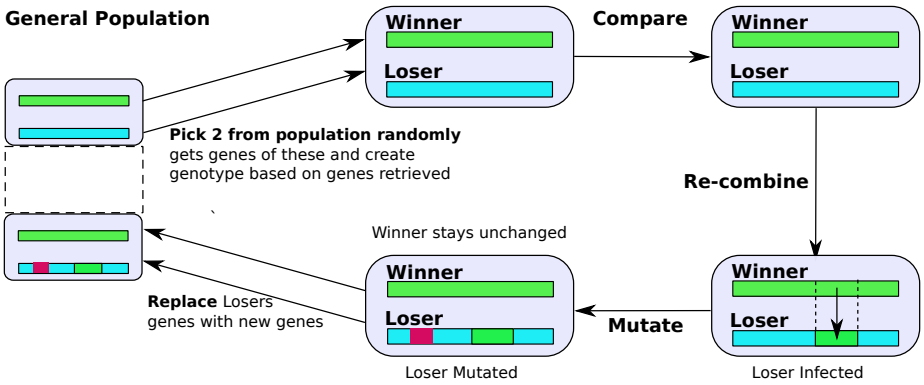

Figure 1: The genotypes of the population are viewed as a pool of strings. One single cycle of the Microbial GA is operated by initially **randomly picking** two, and subsequently **compare** their fitnesses to determine **Winner**, **Loser**, and finally **recombine** where some proportion of Winner's genetic material infects the **Loser**, before **mutating** the revised version of **Loser** (from (Harvey, 2011)).

of parameters to be used and to be updated for subsequent learning. Pathways through the neural network for replication and mutation were selected by a tournament selection genetic algorithm proposed by (Harvey, 2011) during learning. The parameters along an optimal path evolved on the source task were fixed and a new population of paths for the destination task is re-evolved. Such a manner of transfer learning enables the destination task to be learned faster than learning from scratch or fine-tuning, and therefore has greatly succeeded in several supervised learning classification problems (Fernando et al., 2017).

Motivated by the success of PathNet in those applications, in this paper, we explore utilizing PathNet for the facial expression based and speech based emotion recognition tasks. In this work, we first investigate how effective Path-Net is in transferring emotional knowledge from one visual emotion domain to another visual emotion domain to improve overall performance. Based on the experience we have gained in knowledge transferring ability of PahtNet within the visual domain in our previous work (Nguyen et al., 2018), we next investigate whether similar techniques can be used for transferring knowledge within the speech domain. Specifically we investigate the use of PathNet for speech emotion recognition (i) by exploring how well emotional knowledge learned from one speech emotion dataset could be transferred into another

speech emotion dataset, and (ii) by examining how well emotional knowledge learned from multiple speech emotion datasets could be transferred to a single speech emotion dataset.

The contributions of our paper are as follows:

- We introduce a novel transfer learning approach for the emotion recognition task by utilizing PathNet to deal with the problem of insufficient annotated emotion data as well to deal with the catastrophic forgetting issue commonly experienced with traditional transfer learning techniques.

- We confirm, through experimental results, that our proposed system has a significant potential to accurately detect emotions and demonstrates its substantial success in transferring learned knowledge between different emotion datasets, as well as in transferring learned emotional knowledge from multiple speech emotion datasets to a single speech emotion dataset.

- We conduct various sets of within-corpus on three commonly used bench-marking emotion datasets EMODB, eNTERFACE, and SAVEE and we show that the performance of our proposed transfer learning approach for emotion recognition exceeds recent state-of-the-art transfer learning schemes based on fine-tuning/pre-trained models.

The remainder of this paper is organized as follows: Section 2 describes related research; Section 3 presents our proposed system; Section 4 reports our experimental results; and Section 5 concludes the paper.

2. Related work

A number of studies using transfer learning approaches have recently been proposed for the facial expression and speech emotion recognition task. Since the literature review for the facial expression recognition task utilizing deep learning and transfer learning method was reviewed and discussed in our previous work, the interested reader is referred to that work for detailed discussion and analysis, this section will only focus on reviewing the speech emotion recognition task.

2.1. Speech Emotion Recognition using Deep Learning Techniques

Deep learning techniques have emerged as powerful solutions in a wide variety of applications including natural language processing and computer vision owing to their inherent capability of directly learning a hierarchical feature representation from the input data (LeCun et al., 2015). Inspired by their success in multiple fields, many deep learning approaches have been investigated for the task of speech emotion recognition. (Kim et al., 2017a) have proposed a novel architecture for the speech emotion recognition task, in which long short-term memory (LSTM), fully convolutional neural network (FCN), and convolutional neural network (CNN) were combined aiming at extracting local invariant features from the spectral domain. By this combination, long-term dependencies have been well captured, thereby making utterance-level features more discriminative (Kim et al., 2017a). Moreover, by embedding identity skip-connection techniques in their temporal architecture, this proposed system avoided the over-fitting problem caused by training on small amounts of data (Kim et al., 2017a). In another study, in order to handle the large mismatch between training and testing data, (Kim et al., 2017b) have proposed utilizing multi-task learning, and then investigated gender and naturalness as auxiliary tasks of their proposed system. This can enhance significantly the capabilities of generalizing the speech emotion recognition models. The experimental results evaluated on within-corpus and cross-corpus scenarios have shown good performance of their proposed system.

In other approaches, (Nguyen et al., 2017, 2018) have proposed the learning of spatio-temporal features with C3Ds from audio and video for multimodal emotion recognition. (Kim et al., 2017c) have also proposed three dimensional convolutional neural networks (C3Ds) to address the challenge of modeling the spectro-temporal dynamics for speech emotion recognition by simultaneously extracting short-term and long-term spectral features with a moderate number of parameters. (Sahu et al., 2017) have exploited the adversarial auto-encoders focusing on (i) compressing high dimensional vectors encoding emotional utterances into low space vectors (referred to as code vectors) without sacrificing the discrimination during classifying the original vectors, and (ii) generating synthetic samples applying the adversarial auto-encoder, subsequently used for emotion classification. Their system has mainly concentrated on

detecting emotions at utterance level features instead of frame-level ones. (Badshah et al., 2017) have proposed a speech emotion recognition system, in which spectrograms were initially extracted from speech signals at different frequencies. Such spectrograms were subsequently fed into a CNN for emotion prediction. Similarly, (Chang and Scherer, 2017) have addressed emotional valence in human speech by directly learning spectrograms of emotional speech using a CNN. In order to further improve the performance, this architecture has been extended to a deep convolutional generative neural network which was trained in an unsupervised fashion.

Researchers have also explored the whispered speech emotion recognition task, where different feature transfer learning approaches have been explored by utilizing shared-hidden-layer auto-encoders, extreme learning machines auto-encoders, and denoising auto-encoders (Deng et al., 2017a). The key ideas of these approaches were to develop a transformation for automatically capturing useful features hidden in data and to transfer the knowledge from the target domain-testing (whispered speech) to the source domain-training (normal phonated speech), consequently leading the great benefit regarding optimizing all parameters with the support from the test set. Extensive experiments have been conducted with a focus on entirely tackling the binary classification (i.e. valence/arousal) (Deng et al., 2017a). In another study, (Deng et al., 2017b) have also pointed out that many speech emotion recognition systems usually demonstrate poor performance on speech data when there is significant differences between training and test speech arising from the variations in the linguistic content, speaker accents, and domain/environmental conditions (Deng et al., 2017b). To further improve such systems performing under the mismatched training and testing condition, a novel unsupervised domain adaptation algorithm has been introduced and trained by simultaneously learning discriminative information from labeled data and incorporating the prior knowledge from unlabeled data into the learning.

2.2. Speech Emotion Recognition using Transfer Learning Techniques

However, recent studies into emotion recognition have been hindered by the lack of large databases for learning (Zhang et al., 2017; Kaya et al., 2017). To address the lack of large emotion datasets, the fine-tuning/pre-trained model has been recently widely

investigated for the emotion recognition task (Kaya et al., 2017; Zhang et al., 2018; Kim et al., 2017a,c,b; Latif et al., 2018) in which the CNN architectures were pre-trained using the generic ImageNet dataset and fine-tuned on emotion datasets (Ng et al., 2015; Kaya et al., 2017). To learn audio features, (Zhang et al., 2017) used a pre-trained C3D models on large-scale image and video classification datasets, and then fine-tuned them on emotion recognition tasks. To improve the performance of a speech emotion recognition system on such challenging conditions as cross-corpus and cross-language scenario, (Latif et al., 2018) have proposed a transfer learning technique using deep belief networks (DBNs). The experimental results evaluated on five different corpora in three different languages demonstrated the robustness of their system. These results also indicated that use of a large number of languages and a small part of the target data during training could dramatically strengthen the emotion recognition accuracy. However, since they have attempted to validate the system on five different datasets annotated differently, their system only focused on addressing the classification for binary positive/negative valence.

In more recent studies, (Zhang et al., 2018) have reconfirmed that the low-level hand-engineered features seem not to be discriminative enough to recognize the subjective emotions (Zhang et al., 2018). To address this disadvantage, (Zhang et al., 2018) have proposed to extract three channels of log Mel-spectrograms (static, delta, and delta delta corresponding to red, green, and blue in the RGB model of images) from segments over all utterances, and the pre-trained AlexNet model was then fine-tuned on those extracted features. A discriminant temporal pyramid matching technique was subsequently combined with optimal Lp-norm pooling, before exploiting a linear support vector machine to classify the final speech emotion score. Although this architecture demonstrated sufficient performance on tackling discrete speech emotion recognition, the system was unable to address the continuous emotion recognition task (Zhang et al., 2018). (Nguyen et al., 2020) have proposed a joint deep cross-domain transfer learning for emotion recognition which was able to effectively jointly transfer the knowledge learned from rich datasets to source-poor datasets. Moreover, as discussed earlier, all of these fine-tuning approach based systems (Zhang et al., 2018, 2017; Ng et al., 2015; Kaya et al., 2017) still have the drawbacks previously alluded to

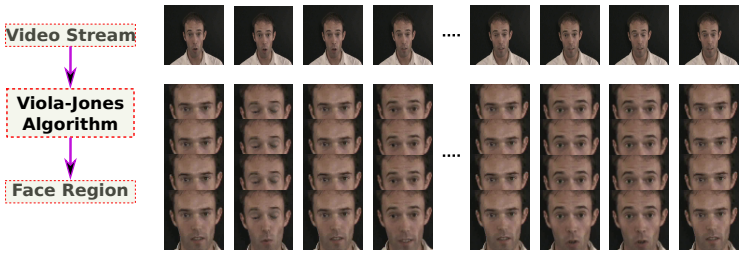

(a) Video pre-processing steps. All frames are initially extracted from all videos, then face regions are detected by applying an improved Viola-Jones algorithm (Nguyen et al., 2018), before being resized to 64×64×3.

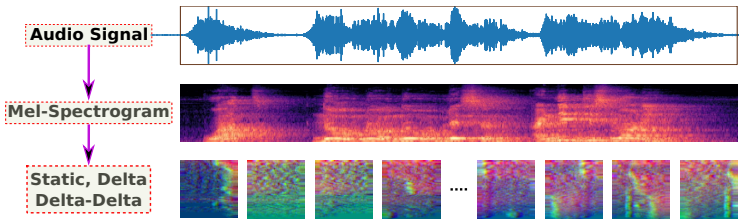

(b) Audio pre-processing steps. First, mel-spectrograms from segments over all utterances are extracted. Then, 3 channels of Mel-spectrograms with size 64 × 64 × 3 (F = 64, T = 64, C = 3) corresponding to the static, delta, and delta-delta coefficients are extracted

Figure 2: This figure illustrates the pre-processing steps for video and audio stream.

such as discarding previously learned information which were detailed by (Rusu et al., 2016). To mitigate these drawbacks, (Gideon et al., 2017) have introduced a learning algorithm using the progressive networks proposed by (Rusu et al., 2016) to effectively transfer knowledge captured from one emotion dataset into another. Although they have handled somewhat successfully the above mentioned limitations, their expensive computation, which kept on increasing when adding new tasks to be learned, makes them less applicable in the implementation of emotion recognition for real world applications.

3. Proposed Methodology

Our main goal in this paper is to investigate techniques to improve the accuracy of deep learning based emotion recognition task which is hindered by the non avail-

ability of large annotated emotion datasets. We achieve this goal by using the recently proposed innovation known as PathNet for transfer learning. We propose the use of a novel transfer learning approach by adopting PathNet to solve the facial expression emotion recognition, and the speech emotion recognition task. Our proposed system is illustrated by simple block diagram consisting of two main blocks which are an input pre-processing block (for video and speech) followed by our proposed PathNet block along with the output classifying block (as illustrated in Fig. 3). The video and audio stream are initially pre-processed. For video stream, we initially exploit a Viola Jones-based algorithm (Nguyen et al., 2017) to extract all face regions from both SAVEE and eNTERFACE (Martin et al., 2006) datasets. This is described in detail in Section 3.1. For audio stream, we also initially extract three channels of log Mel-spectrograms (static, delta, and delta delta corresponding to red, green, and blue) from segments over all utterances from eNTERFACE (Martin et al., 2006), SAVEE (Haq and Jackson, 2010), and EMO-DB (Burkhardt et al., 2005) and these steps are described in detail in Section 3.2. These extracted video and audio features are subsequently fed into our PathNet to classify a final facial expression and speech emotion score, respectively. To the best of our knowledge, the use of PathNet has not been previously investigated in dealing with the dearth of suitable emotion databases for the development of emotion recognition system. The procedures of feature extraction and our PathNet architecture are described in more detail in the following subsections.

3.1. Video pre-processing

All frames are initially extracted from visual signal for further steps. Since such extracted frames still contain considerable redundant information for emotion detection, we extract only the face regions using the simple algorithm (Nguyen et al., 2017) as follows:

1. All bounding boxes containing face regions in each frame were extracted employing the Viola-Jones algorithm (Viola and Jones, 2004) and a face region was then detected.
2. In some cases where the Viola Jones algorithm detected no faces, or more than 1 face, the location of the previously detected face region was used.

By applying this algorithm, we have successfully extracted all face regions from all frames in both datasets (SAVEE and eNTERFACE) as input into PathNet (see in Fig. 2 (a). as an illustration of some input samples)

3.2. Audio pre-processing

(Zhang et al., 2018) have pointed out that the low-level hand-engineered features, which includes RASTA-PLP (Hermansky et al., 1992), pitch frequency features, energy-related features (Ververidis et al., 2004), formant frequency (Xiao et al., 2005), Zero Crossing Rate (ZCR) (Rabiner and Sambur, 1975), Mel-Frequency Cepstrum Coefficients (MFCC) and its first derivative, Linear Prediction Cepstrum Coefficients (LPCC), Linear Prediction Coefficients (LPC) (Pao et al., 2006; Shen et al., 2011), seem not to be discriminative enough for recognizing the subjective emotions (Zhang et al., 2018). Therefore, in order to boost the performance of speech emotion recognition system, instead of exploiting such hand-crafted features, (Zhang et al., 2018) have extracted three channels of log Mel-spectrograms from segments over all utterances, and then fine-tuned the pre-trained AlexNet model on these extracted features. Their proposed architecture has successfully addressed the speech emotion recognition task. Motivated by this high performance, in this paper we also propose to extract these features as follows (see Fig. 2 (b)):

- Mel-spectrogram segments with size 64 $\times$ 64 $\times$ 3 (F = 64, T = 64, C = 3) are generated from 1-D speech signals. F denotes as the number of Mel-filter banks, T denotes as the segment length corresponding to the frame number in a context window, and C represents the number of channels of Mel-spectrogram. The 3 channels of Mel-spectrograms are the static, delta and, delta-delta coefficients of Mel-spectrogram), respectively (Zhang et al., 2018).

- Specifically, 64 Mel-filter banks from 20 to 8000 Hz are exploited to extract the whole log Mel-spectrogram using a 25ms Hamming window size with 10ms overlapping for an utterance. A context window of 64 frames (its length 10ms $\times$ 63 + 25ms = 655ms) is then applied to the whole log Mel-spectrogram to extract

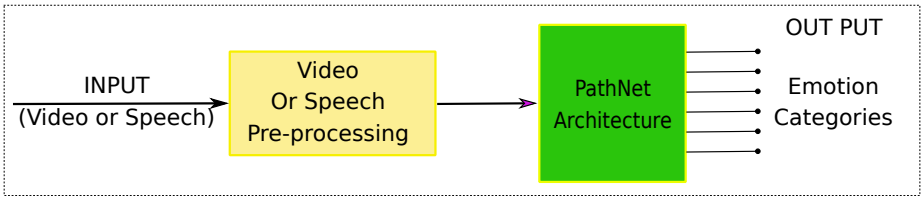

(a) Our proposed emotion recognition block diagram

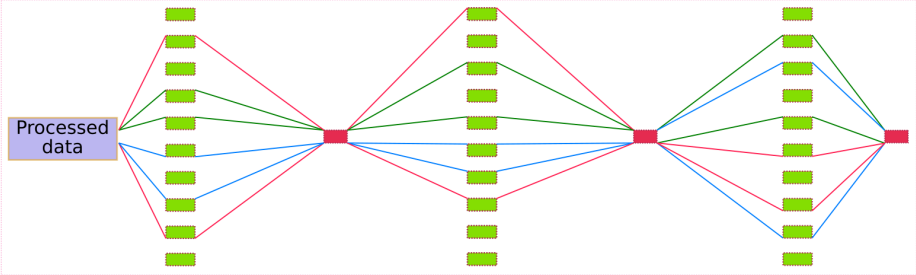

(b) The PathNet architecture has 3 layers, and each layer consists of 20 modules, and each module is followed by a rectified linear unit. The feature maps between layers are averaged before being fed into the modules of the next layer. A maximum of 4 modules per layer are typically allowed to be presented in a pathway.

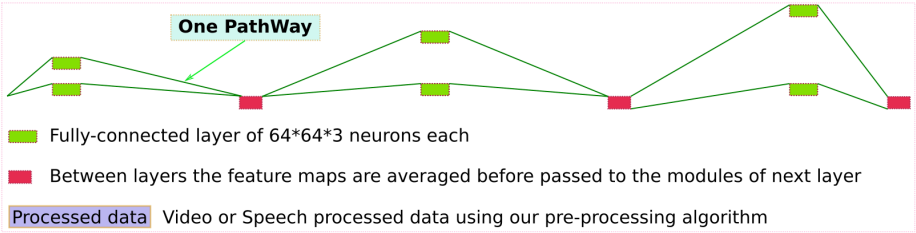

(c) Description of each component of our PathNet architecture

Figure 3: Illustrates our proposed emotion recognition system

the static 2-D Mel-spectrogram segments (64 $\times$ 64) with overlapping size of 30 frames (Zhang et al., 2018).

3.3. PathNets

Our PathNet architecture and its settings relies on PathNets (Fernando et al., 2017) which was used to conduct all sets of experiments on CIFAR (Krizhevsky and Hinton, 2009) and SVHN (Netzer et al., 2011). The following sections will provide in detail

its architecture and explain further how to train our system, how to transfer learned emotional knowledge between emotion dataset.

3.3.1. PathNet architecture

Our PathNets includes a number of layers (L = 3), a number of modules (M = 20) per layer of 20 neurons in each. Each module itself functions as a neural network consisting of linear units, and followed by a transfer function (rectified linear units adopted). For each layer the outputs of the modules of this layer are averaged before being fed into the active modules of the subsequent layer. A module is active if it is shown in the path genotype and currently validated (shown in Fig. 3 (b)). A maximum of 4 distinct modules per layer are typically allowed in a pathway. The final layer is not shared for each task which is being learned (Fernando et al., 2017).

3.3.2. Pathway Evolution and Transfer Learning Approach

One emotion dataset (source data) is trained for a fixed number of generations with a goal of finding an optimal pathway by adopting a binary tournament selection algorithm proposed by (Harvey, 2011) (see Fig. 1) which takes responsibility for eliminating bad configurations and mutating good ones, and subsequently training them further. This pathway is then fixed. This means that its parameters are no longer permitted to modify and the rest of parameters, which are not shown in such best fit path, are reinitialized, and are then again trained/evolved on the another emotion dataset (destination data). Through this knowledge transferring approach, the destination data is permitted to be learned faster than learning from scratch or after fine-tuning. The performance measurement of our proposed system is the recognition accuracy achieved after such fixed training time. Evidence to confirm a positive transfer in these cases is given by a better final recognition accuracy of our proposed system achieved when trained on destination data than that achieved by learning from scratch.

When PathNet is trained on source data, at the beginning, a population of genotypes is randomly generated (See Fig. 4). In each generation, two pathways are randomly selected to train on source data (see Fig. 5). The reason only two pathways are selected is that binary tournament selection is exploited to choose the pathways. One

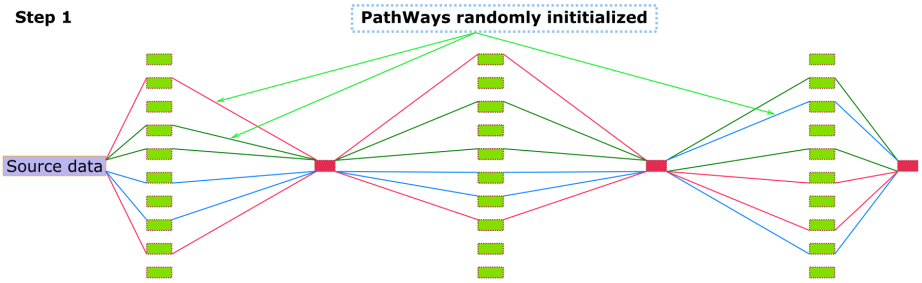

Figure 4: A number of pathways is randomly initialized when pathnet learning on source data

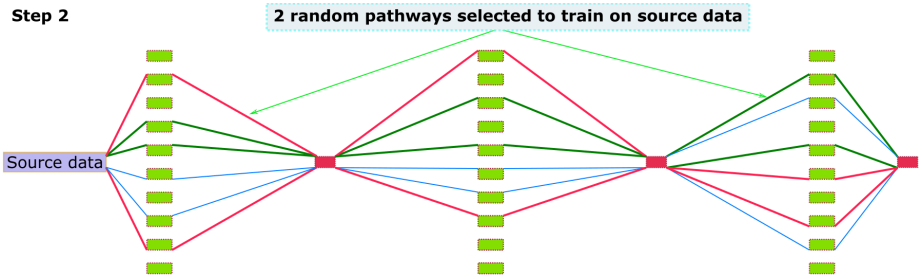

Figure 5: Two pathways are selected to train on source data

pathway is learned using stochastic gradient descent for T epochs (T equals the number of samples counted in training set divided by mini-batch size). Another pathway is also learned using stochastic gradient descent for T epochs (T equals the number of samples counted in training set divided by mini-batch size). The fitness of such pathways is the rate of correct samples on the training set during that period of training time. When training two pathways is completed, the pathway with the bad performance (called loser) is replaced by the pathway with the better performance (called winner) and the winner is then mutated with equal probability $1/(4\times3)$ per each candidate of the genotype (see Fig. 6, and Fig. 7), a new random integer from range [-2, 2] is added to the current value of the winner candidate. We repeat step 2 (Fig. 5) and step 3 (Fig. 6) for number of generations. When training on source data is completed we achieve best pathway (see Fig. 8).

The parameters presented in this best fit pathway are fixed and are reused for training on the destination data and the rest of parameters are randomly reinitialized. When training PathNet on the destination data, the procedures for this training stage are sim-

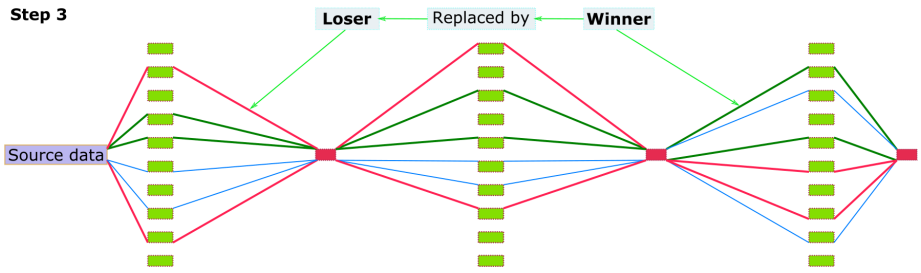

Figure 6: A pathway with bad performance (loser) is replaced by a pathway with better performance (winner)

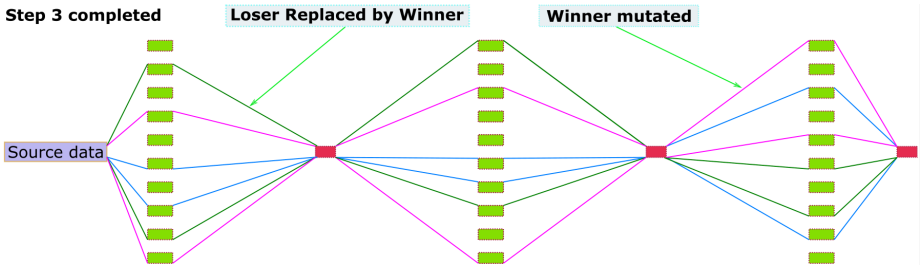

Figure 7: Winner is mutated

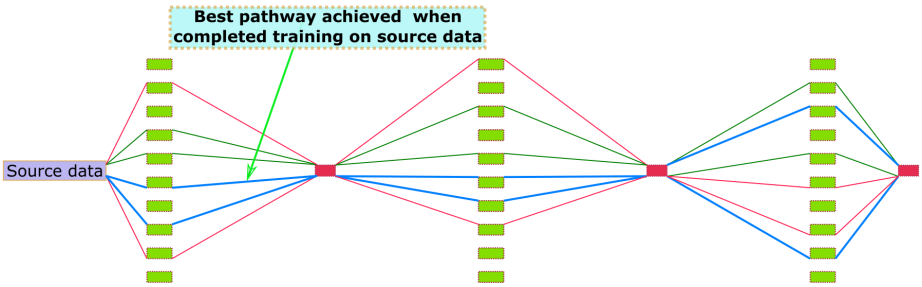

Figure 8: Best pathway achieved when completed training on the source data

ilar to ones when training on the source data (procedures illustrated in Fig. 9, Fig. 10, and Fig. 11). At the beginning, a new population of genotypes is randomly initialized and then trained/evolved further on the destination data. However, the difference is that the best pathway achieved when completed training on source data is always activated during learning with the destination data. When completed training on the destination data, we also achieve best pathway (see Fig. 13).

It is quite difficult to fully understand how to train/test PathNet and how it works

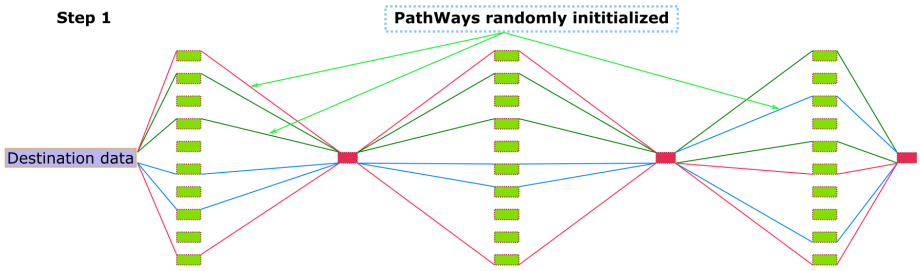

Figure 9: A number of pathways are randomly initialized when pathnet trained on the destination data

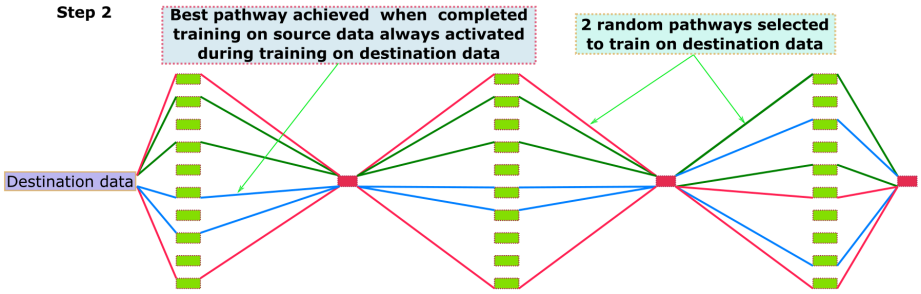

Figure 10: Two random pathways are also selected to train on the destination data

by only reading textual explanation that is presented in this section. Therefore, in order to make our PathNet based transfer learning approach more understandable, we have added further explanation on the progression of achieving the best pathway by visualizing every step using the corresponding figure (please see Figs. 3.3 - 3.13). As we can see these Figs, to ease for readers to view and follow step by step, we simplify the architecture by drawing only ten modules in each layers, up to two modules are activated in each layer, and a population of four random pathways are initialized. Whereas the Figs. 3.19- 3.23 are drawn to further explain and visualize the progression of how to achieve best pathway for one particular set of experiment. These Figs illustrate exact parameters and architectures corresponding to the progression in gaining the optimal pathway when conducting the set of experiment. We believe that this explanation is more comprehensive than done by the original PathNet paper (Fernando et al., 2017) as in that paper authors did not illustrate such steps using figures.

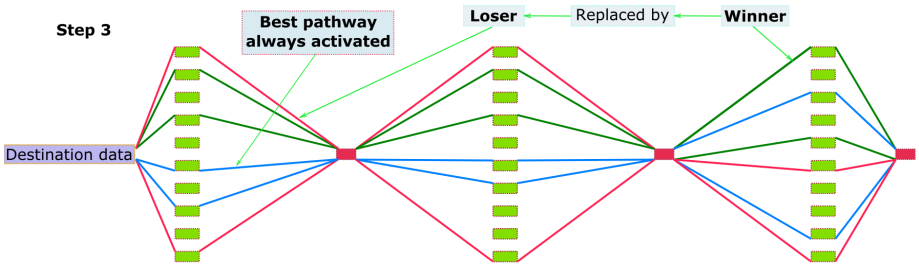

Figure 11: A pathway with bad performance (loser) is replaced by a pathway with better performance (winner) when learning with the destination data

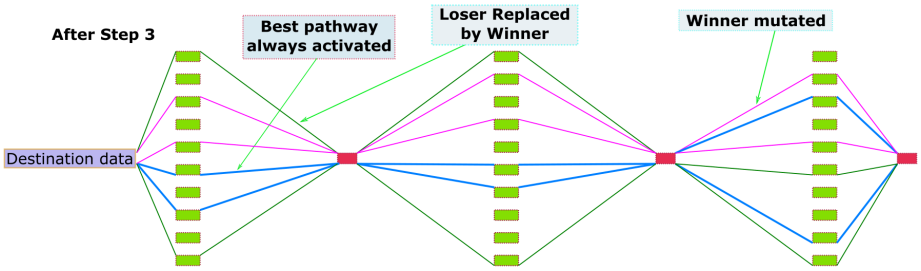

Figure 12: The winner is then mutated

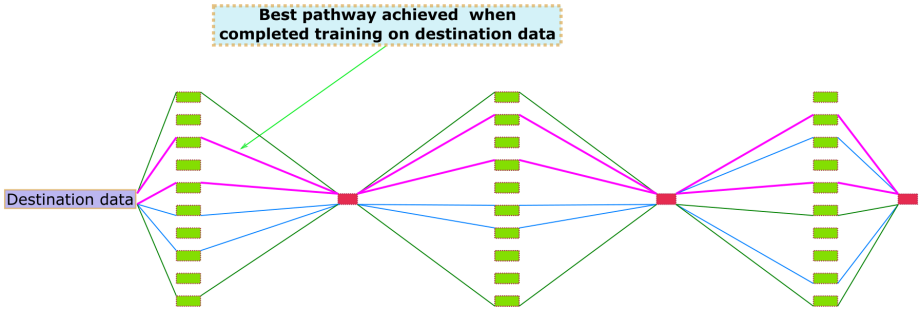

Figure 13: Best pathway is achieved when completed training on the destination data

4. Experiments & Results

Dataset Details: The eNTERFACE dataset (Martin et al., 2006) is an audio-visual dataset which has 44 subjects and includes a total of 1293 video sequences in which the proportion of sequences corresponding to women and men are 23% and 77%, respectively. They were asked to express 6 discrete emotions including anger, disgust, fear, happiness, sadness, and surprise (Martin et al., 2006).

The SAVEE dataset (Haq and Jackson, 2010) is an audio-visual dataset which was recorded by higher degree researchers (aged from 27 to 31 years) at the University of Surrey, and four native male British speakers. All of them were also required to speak and express seven discrete emotions such as anger, disgust, fear, happiness, sadness, surprise, and neutral. The dataset comprises of 120 utterances per speaker, resulting in a total of 480 sentences (Haq and Jackson, 2010).

The EMO-DB dataset (Burkhardt et al., 2005) is an acted speech corpus containing 535 emotional utterances with seven different acted emotions listed as disgust, anger, neutral, sadness, boredom, and fear. These emotions were stimulated by five male and five female professional native German-speaking actors, generating five long and five short sentences German utterances used in daily communication. These actors were asked to read predefined sentences in the targeted seven emotions. The audio files are on average around 3 seconds long were recorded using an anechoic chamber with high-quality recording equipment at a sampling rate of 16 kHz with a 16-bit resolution and mono channel (Zhang et al., 2018).

Performance Measures: We evaluate our proposed system on all three datasets: the SAVEE, the eNTERFACE, and the EMO-DB dataset. We apply k-fold cross-validation, the original training data is randomly divided into k equal parts. Of the k-parts, one of them is fixed as the validation data for testing the model, and the other k-1 parts are used as training data. The cross-validation process is then repeated 5 times. We also apply the leave-one-subject-out cross-validation protocol which means that we have to conduct N experiments if the dataset consists of N subjects. For each experiment, N-1 subjects are included in the training set and the remaining subject is used for the testing set and all of our sets of experiments are carried out in a subject independent manner.

Apart from these, Weighted Averaged Recall (WAR) has been recently adopted and has become widely a standard measure for evaluating the performance of emotion recognition systems (Zhang et al., 2018). In addition, Unweighted Averaged Recall (UAR) (Eyben et al., 2016; Eyben, 2016; Zhang et al., 2018) has been also popularly adopted to evaluate this performance with respect to reflecting unbalance between emotional classes (Zhang et al., 2018). Therefore, in order to fairly compare

Table 1: Different measures for multi-class classification C_i. For class C_i, tp_i are true positive, and fp_i - false positive, fn_i - false negative, and tn_i - true negative counts, l - the number of classes, (Sokolova and Lapalme, 2009).

Measure	**Formula**	**Evaluation focus**
Unweighted Averaged Precision (UAP)	$\frac{\sum_{i=1}^{l} \frac{tp_i}{tp_i+fp_i}}{l}$	An average per-class agreement of the data class labels with those of a classifiers
Unweighted Averaged Recall (UAR)	$\frac{\sum_{i=1}^{l} \frac{tp_i}{tp_i+fp_i}}{l}$	An average per-class effectiveness of a classifier to identify class labels

with some recent state-of-the-art speech emotion recognition schemes using the same above-mentioned measures (WAR, UAR), in this paper, we also compute and compare both Weighted Averaged Recall and Unweighted Averaged Recall to validate the performance of our proposed system. Table 1 illustrates the detail formulas on how to calculate these evaluating measures.

As our baseline systems, we use the methodology proposed by (Zhang et al., 2018) and a off-the-shelf CNN (AlexNet) as our baseline systems. Additionally, we also implement additional baseline systems and each of which is trained on the same emotion dataset from scratch using PathNet.

Since our main focus of this paper is to address the issue of lack of emotion data for deep learning techniques. We clearly introduced this problem in the **Abstract** and further discussed it in the **Introduction** section. To make compatible with the problem we are trying to solve, we should choose small emotion datasets to show that our proposed methodology demonstrates a robust performance on such poor data. Furthermore, all large-scale datasets such as CK+, RAF-DB and AffectNet only can be used for facial expression recognition which is not the main task of this paper. It is, however, noteworthy that addressing the issue of insufficient data for speech emotion recognition task is the main work of our paper.

We conduct various sets of experiments using our proposed system in examining

Table 2: Description of notation of all models for all sets of experiments of our proposed system conducted on eNTERFACE, SAVEE, and EMODB;

Transferring emotional knowledge	**Notation**
PathNet is trained from scratch on visual eNTERFACE	V_eNTER
PathNet is trained from scratch on visual SAVEE	V_SAV
PathNet is trained from scratch on audio SAVEE	A_SAV
PathNet is trained from scratch on audio EMODB	A_EMO
From visual eNTERFACE to visual SAVEE	V_eNTER→ SAV
From visual SAVEE to visual eNTERFACE	V_SAV→ eNTER
From audio eNTERFACE to audio SAVEE	A_eNTER→ SAV
From audio eNTERFACE and audio SAVEE to audio EMODB	A_eNTER+SAV→ EMO

Table 3: Results of our proposed system evaluated on visual SAVEE, visual eNTERFACE, audio SAVEE, and audio EMODB. In the first two rows, we explore transferring the emotional knowledge from visual eNTERFACE to visual SAVEE and vice versa. In the last two rows, the emotional knowledge is transferred from audio eNTERFACE to audio SAVEE. We also explore transferring the emotional knowledge from multiple audio emotion domains (audio eNTERFACE and audio SAVEE) to one audio emotion domain (audio EMODB).

Method	**Ang**	**Sur**	**Dis**	**Fea**	**Hap**	**Sad**	**WAR**
$\mathbf{V_{eNTER \rightarrow SAV}}$	0.93	0.94	0.96	0.92	0.94	0.95	0.94
$\mathbf{V_{SAV \rightarrow eNTER}}$	0.99	1.0	0.95	0.81	1.0	0.93	0.94
$\mathbf{A_{eNTER \rightarrow SAV}}$	0.73	0.70	0.56	0.72	0.80	0.77	0.71
$\mathbf{A_{eNTER+SAV \rightarrow EMO}}$	0.99	0.99	1.0	0.94	0.94	0.96	0.97

how well the accuracy of facial expression and speech emotion recognition are improved when transferring learned emotional knowledge between emotion datasets. In the first two sets of experiments, we conduct experiments for facial expression recognition, in which we transfer emotional knowledge from visual emotion eNTERFACE dataset (source data) to visual emotion SAVEE dataset (destination data) and vice versa. For the next two sets of experiments, we conduct experiments for speech emotion recognition by transferring emotional knowledge from audio emotion eNTERFACE dataset (source data) to audio emotion SAVEE dataset (destination data), and transferring emotional knowledge from multiple audio emotion datasets (eNTERFACE and SAVEE) to audio emotion EMODB.

PathNet Settings: Our proposed PathNet architecture (i.e. $L = 3$ layers, M = 20 liner units per layer of 20 neurons each followed by rectified linear units, average function used to activate units between two layers, and a maximum of 4 of those units per layer included in a pathway which is represented by a 4×3 matrix of intergers in the range [1,13]). Our proposed PathNet architecture is trained on visual eNTERFACE/visual SAVEE/audio eNTERFACE/(audio eNTERFACE and audio SAVEE) which are set as as source data and is then trained/evolved on visual SAVEE/visual eNTERFACE/audio SAVEE/audio EMODB, respectively. In both source tasks and destination tasks of these sets of experiments, our proposed systems are trained for 200 generations. At the beginning of each task, a population of 20 of pathways are randomly generated and are then trained on the source data. In each generation, two paths are randomly selected to train on the source data and then on the destination data for validation. To evaluate one pathway, a pathway is trained with stochastic gradient descent with learning rate 0.02, a mini-batch size of 64 and is trained for T epochs (T equals the number of samples counted in training set divided by mini-batch size of 64). The fitness of such pathway is the rate of correct samples on the training set during that period of training time. When completing the calculation of the fitness of two pathways, the pathway with the smaller fitness is replaced by the pathway with the greater one that is then mutated with equal probability $1/(4\times3)$ per each candidate of the genotype, a new random integer from range [-2, 2] is added to the current value of the lose candidate. Therefore, the network between tasks is modified

Table 4: Results of our proposed system when transferring the emotional knowledge from visual eNTERFACE to visual SAVEE in comparison with the best baseline system.

Method	**WAR (%)**
Fine-Tuned AlexNet (Zhang et al., 2018)	0.85
V_{SAV}	0.89
$\mathbf{V_{eNTER \rightarrow SAV}}$	**0.94**

Table 5: Results of our proposed system when transferring the emotional knowledge from visual SAVEE to visual eNTERFACE in comparison with the best baseline system.

Method	**WAR (%)**
Fine-Tuned Alexnet (Zhang et al., 2018)	0.88
V_{eNTER}	0.88
$\mathbf{V_{SAV \rightarrow eNTER}}$	**0.94**

as follows: the parameters presented in the best fit pathway, which is evolved on the source data, are fixed and the rest of parameters are randomly reinitialized and a new population of 20 of genotypes are randomly initialized and then are trained/evolved further on the destination data.

4.1. Experimental Results

In this section, we report, analyze, and compare experimental results of our proposed system on aforementioned sets of experiments. In all the experiments, we have taken meticulous care to ensure that the test data is never used for training.

In the first two sets of experiments, we conduct experiments to show the robustness of our proposed system when performing under the condition of insufficient data for facial expression recognition. Our proposed system is initially trained on visual eNTERFACE, and is then evolved on visual SAVEE and vice versa. Since SAVEE dataset

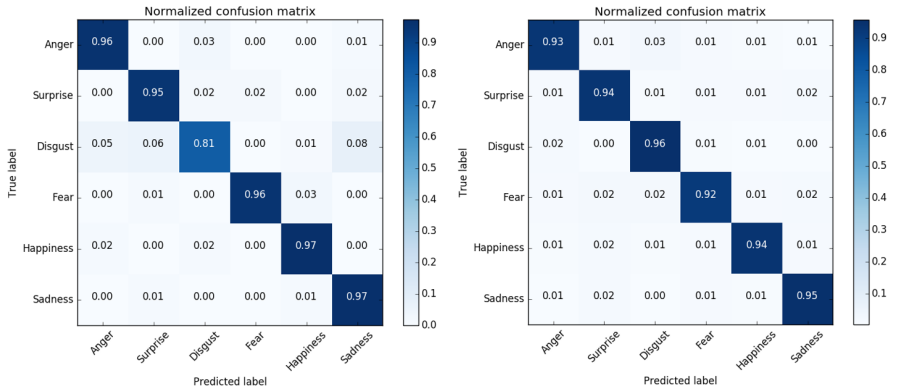

(a) Confusion matrix of our proposed system when transferring the emotional knowledge from visual eNTERFACE to visual SAVEE

(b) Confusion matrix of our proposed system when transferring the emotional knowledge from visual SAVEE to visual eNTERFACE

Figure 14: Illustrates the confusion matrix of our proposed system evaluated on visual SAVEE and visual eNTERFACE

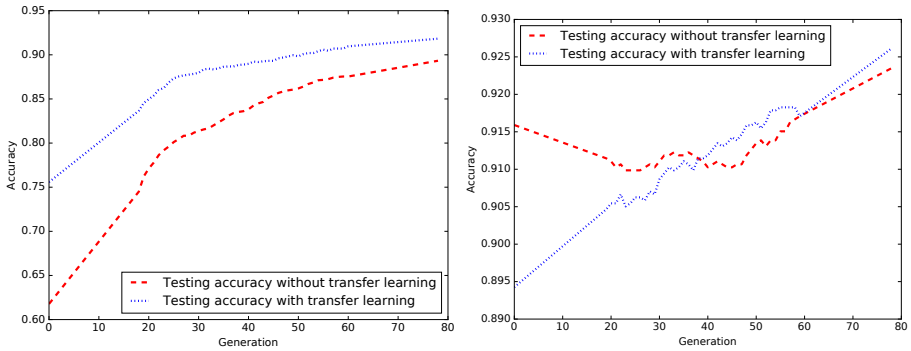

(a) Testing learning curves on visual SAVEE

(b) Testing learning curves on visual eNTERFACE

Figure 15: Illustrates testing learning curves of our proposed system during learning on visual SAVEE and visual eNTERFACE with and without transfer learning.

consists of an additional type of emotion (neutral) compared to eNTERFACE dataset, hence to be consistent between these two datasets, only shared types of emotion including anger, surprise, disgust, fear, happiness, and sadness are detected via these two sets of experiments. The experimental results of our proposed system on this evaluation are shown in the first two rows of Table 3.

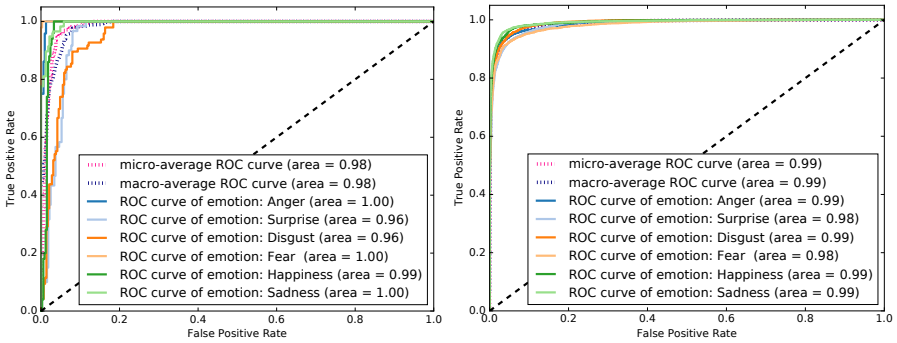

(a) Receiver operating characteristic (ROC) curves of individual emotions when transferring the emotional knowledge from visual eNTERFACE to visual SAVEE

(b) Receiver operating characteristic (ROC) curves of individual emotions when transferring the emotional knowledge from visual SAVEE to visual eNTERFACE

Figure 16: Receiver operating characteristic (ROC) curves of our proposed system on visual SAVEE and visual eNTERFACE

Visual eNTERFACE $\rightarrow$ Visual SAVEE: As illustrated in Table 3, our proposed system ($V_{eNTER \rightarrow SAV}$), which transfers the emotional knowledge from visual eNTERFACE to visual SAVEE, achieves 94% of facial expression recognition accuracy in regard to WAR. This accuracy is 5% and 9% significant higher than those achieved by our facial expression recognition baseline systems: V_{SAV} which is trained on visual SAVEE from scratch and the system fine-tuning on AlexNet proposed by (Zhang et al., 2018), respectively (see Table 4).

Visual SAVEE $\rightarrow$ Visual eNTERFACE: To further show the efficiency of our proposed system in solving the issue of insufficient facial expression data, we explore transferring the emotional knowledge from visual SAVEE to visual eNTERFACE. The experimental results are illustrated in the second row of Table 3. In this setting, our proposed system also performs significantly better than our facial expression recognition baseline systems, achieving the best facial expression recognition accuracy regarding to WAR (94%), which is 6.5% and 6.35% better than those obtained by our facial expression recognition baseline systems: V_{eNTER} which is trained on visual eNTER from scratch and the system fine-tuning on AlexNet proposed by (Zhang et al., 2018)

Table 6: Results of our proposed system when transferring the emotional knowledge from audio eNTERFACE to audio SAVEE in comparison with the best baseline speech emotion recognition system.

Method	**WAR**
AlexNet (Zhang et al., 2018)	0.69
A_{SAV}	0.81
$\mathbf{A_{eNTER \rightarrow SAV}}$	0.85

(illustrated in Table 5).

To depict the performance of individual emotion of these two systems, the confusion matrices of their systems ($V_{eNTER \rightarrow SAV}$ and $V_{SAV \rightarrow eNTER}$) are illustrated in Fig. 14 (a), and Fig. 14 (b), respectively. We also visualize the robust performances of each emotion of these two systems, and plot their ROC curves which are illustrated in Fig. 16 (a) and Fig. 16 (b), respectively. Moreover, in order to further visualize the effectiveness of our proposed systems , the different testing learning curves are plotted as illustrated in Fig. 15. As demonstrated in Fig. 15, the performance of both proposed systems with transfer learning surpasses considerably those without transfer learning.

In the first two sets of experiments, we have focused on conducting experiments for facial expression recognition. Through experimental results, we have shown that our proposed facial expression recognition system is vastly boosted when transferring the emotional knowledge from one visual emotion domain to another visual emotion domain using PathNet. The performance of our system is also superior to those of the recent state-of-the-art facial expression recognition systems fine-tuning on the off-the-shelf/pre-trained models. To further demonstrate the effectiveness of our proposed system, we carry out extensive experiments for speech emotion recognition task.

Audio eNTERFACE $\rightarrow$ Audio SAVEE: In this set of experiments, the emotional knowledge is transferred from audio eNTERFACE to audio SAVEE. Experimental results are shown in Table 6. It can be seen that our proposed system ($A_{eNTER \rightarrow SAV}$ with the best speech emotion recognition accuracy of 85%) demonstrates a significant superior performance over the speech emotion recognition system proposed by (Zhang

Table 7: Results of our proposed system when transferring the emotional knowledge from audio eNTERFACE and audio SAVEE to audio EMODB in comparison with the best baseline speech emotion recognition system

Method	**WAR (%)**
Fine-Tuned Alexnet-Average (Zhang et al., 2018)	0.83
A_{EMO}	0.89
$\mathbf{A_{eNTER+SAV \rightarrow EMO}}$	**0.97**

et al., 2018) and our baseline system (A_{SAV} which have been trained on audio SAVEE from scratch using PathNet) by 16% and 4% regarding WAR, respectively.

(Audio eNTERFACE and Audio SAVEE) $\rightarrow$ Audio EMODB: In this set of experiment, our proposed system ($A_{eNTER+SAV \rightarrow EMO}$ with the best speech emotion recognition accuracy of 97%) is trained on audio eNTERFACE and audio SAVEE and the learned emotional knowledge (parameters) presented in the best pathway is then transferred and always involved in training stage on audio EMODB. We apply the leave-one-subject-out cross-validation (LOSOCV) for this set of experiment. The reason we investigate these settings for this set of experiments is that in the first three sets of experiments, we have only explored transferring the emotional knowledge from one emotion domains to another emotion domain, to gain more emotional knowledge, in this set of experiment we explore transferring the emotional knowledge from multiple emotion domains to another emotion domain and use LOSOCV to be compatible with our baseline system. The experimental results are illustrated in Table 7. It can be clearly shown in Table 7 that our proposed system (97%) performs significantly better than the transfer learning approach based on pre-trained/fine-tuning models (Zhang et al., 2018) by 14% and our baseline system (A_{EMO} which have been trained on audio EMODB from scratch using PathNet) by 8% in regard to WAR.

Similarly, in order to further gain insight into the performances of individual speech emotions of both systems ($A_{eNTER \rightarrow SAV}$ and $A_{eNTER+SAV \rightarrow EMO}$), we illustrate the confusion matrix of these speech emotion recognition systems (see Fig. 18 (a), and Fig.

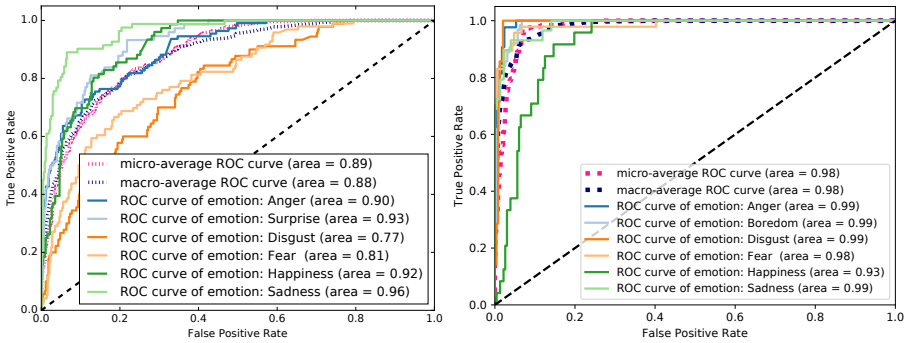

(a) Receiver operating characteristic (ROC) curves of individual emotions when transferring the emotional knowledge from audio eNTERFACE to audio SAVEE

(b) Receiver operating characteristic (ROC) curves of individual emotions when transferring the emotional knowledge from audio eNTERFACE and audio SAVEE to audio EMODB

Figure 17: We illustrate receiver operating characteristic (ROC) curves of individual emotions evaluated on audio SAVEE and audio EMODB

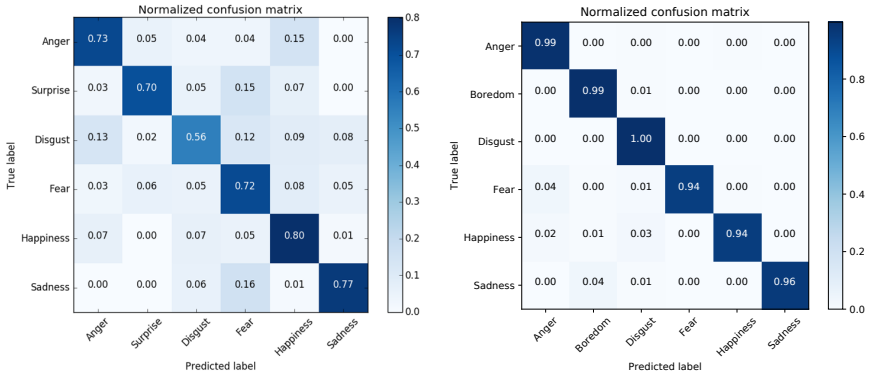

(a) Confusion matrix of our proposed system when transferring the emotional knowledge from audio eNTERFACE to audio SAVEE

(b) Confusion matrix of our proposed system when transferring the emotional knowledge from audio eNTERFACE and audio SAVEE to audio EMODB

Figure 18: Illustrates confusion matrix of our proposed system evaluated on audio SAVEE and audio EMODB

18 (b), accordingly). Moreover, we have also plotted their corresponding receiver operating characteristic (ROC) curves as illustrated in Fig. 17 (a) and Fig. 17 (b). Specif-

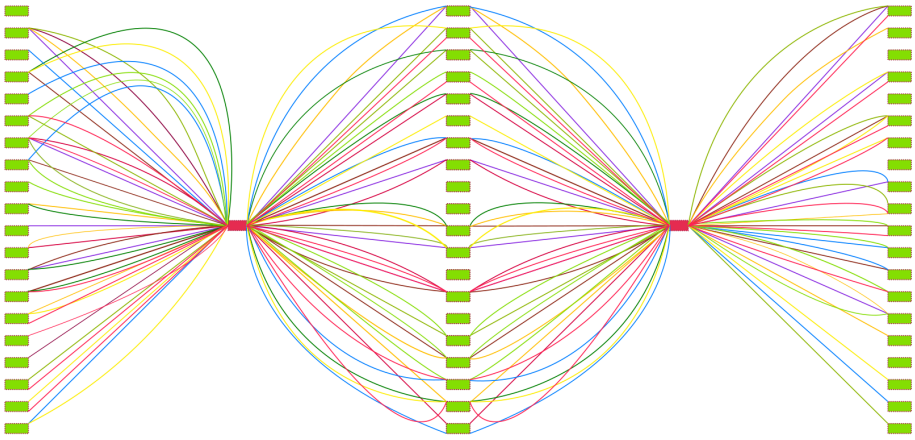

Figure 19: A population of 20 pathways are randomly initialized when learning on the source data (audio eNTERFACE).

ically, our system ($A_{eNTER \rightarrow SAV}$) recognizes sadness best with AUC = 0.96, whereas it does not detects disgust and fear so well (with only AUC = 0.77 and AUC = 0.81, accordingly). For other speech emotions including anger, surprise, and happiness our proposed system also demonstrates high-potential performance (with AUC = 0.90, 0.93, and 0.92, respectively) (see Fig. 17 (a)). It can also be seen that our $A_{eNTER+SAV \rightarrow EMO}$ reveals a great success when detecting speech emotions such as anger, boredom, disgust, fear, and sadness with their corresponding AUC = 0.99, 0.99, 0.99, 0.98, and 0.99, whereas it performs less effectively on the happiness emotion with only AUC = 0.93 (as shown in Fig. 17 (b)).

Discussion: The reason we achieve significantly better performance when transferring the emotional knowledge from one emotion dataset to another emotion dataset using PathNet is that the emotional knowledge presented in the best PathWay achieved when completed training on one emotion dataset (source data) is now reused as apart of initial emotional knowledge and is always involved in training stage on another set of pathways of PathNet on another emotion dataset (destination data). However, when PathNet is trained on the same emotion dataset (destination data) from scratch, the emotional knowledge is randomly initialized and is learned with only one set of pathways on only one emotion dataset (destination data).

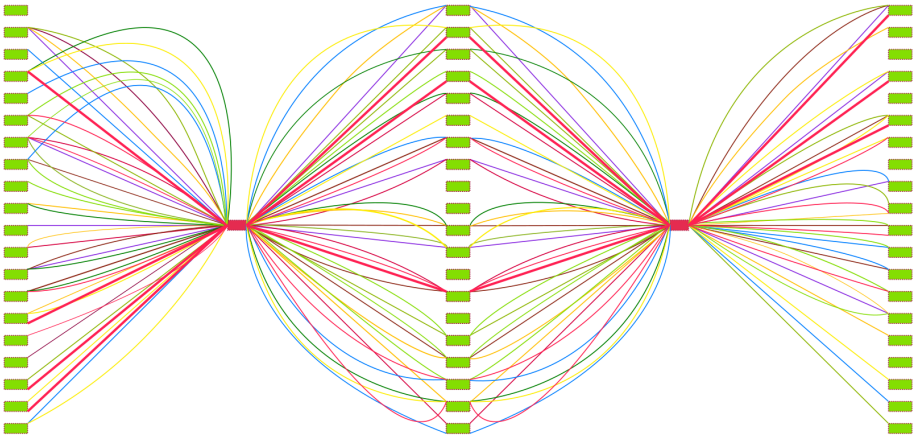

Figure 20: The optimal pathway (highlighted by red colour) is achieved when the training stage is completed on audio eNTERFACE as the source data and the parameters presented in the pathway are fixed.

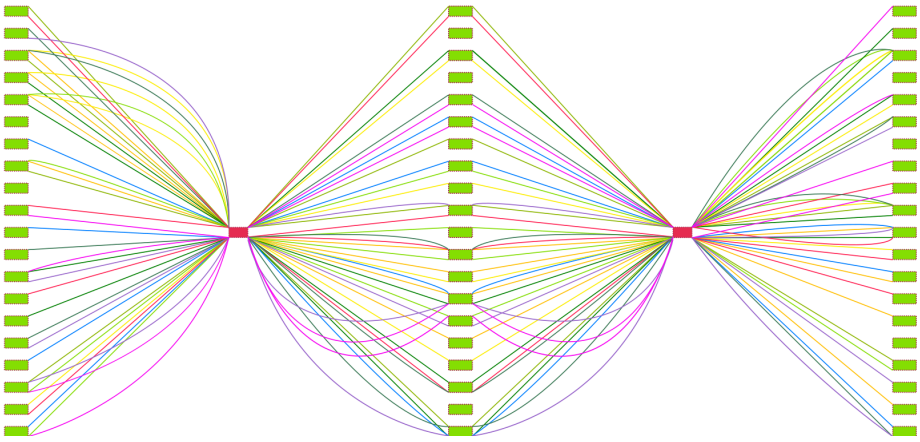

Figure 21: A new population of pathways are generated when learning on the destination data (audio SAVEE)

To the best of our knowledge, fine-tuning is arguably the most widely exploited method for transfer learning while working with deep learning architectures. It begins with a pre-trained model on the source task and further trains it on the target task. ImageNet pre-trained models are commonly used for fine-tuning. Compared with training from scratch, fine-tuning a pre-trained off-the-shelf CNN on a target dataset can considerably boost performance, whereas lessening the target annotated data requirements. This transfer learning approach is also considered as one of the most recent state-of-

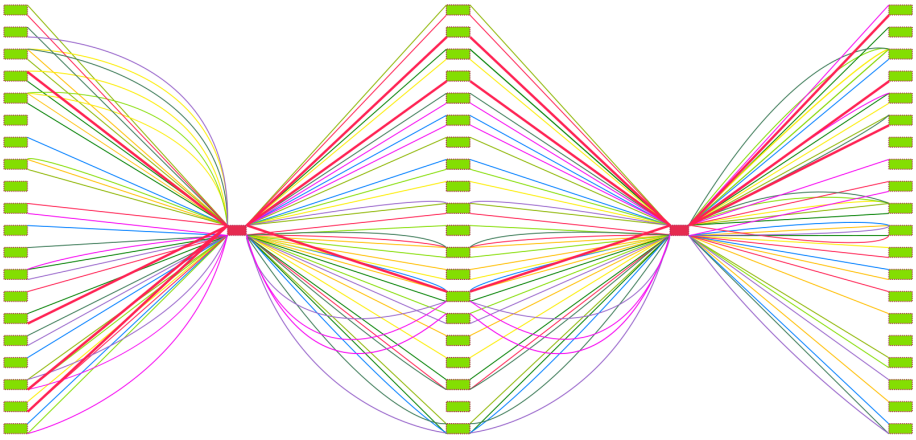

Figure 22: The pathway highlighted by red colour, which have been transferred from the model on the source data, are always activated during evolving on the destination task (audio SAVEE), however, their parameters are fixed.

the-art methods. However, as pointed out earlier pre-trained models are still limited in their generalization capability and thus lead to poor performance on novel test sets. The main objective of this paper is to propose an alternative transfer learning approach to address such disadvantages. Therefore, we have primarily compared our transfer learning method using PathNet with a state-of-the-art transfer learning method relying on fine-tuning/pre-trained models as reported in Experiments & Results Section. The transfer learning method that we have compared against my PathNet based approach using the same datasets is fined-tuned AlexNet. Moreover, unfortunately we could not find any other state of the art transfer learning method which uses the same datasets to make a fair comparison. Our proposed system outperforms the state-of-the-art emotion recognition system relying on the transfer learning method which fine-tunes off-the-shelf/pre-trained models. We believe that this is due to the fact that the representational features that are unrelated to emotion are still retained in off-the-shelf/pre-trained models and the extracted features are also vulnerable to identity variations in these approaches, leading to degrading the performance of the emotion recognition system fine-tuning off-the-shelf/pre-trained models on the emotion dataset.

As presented in Proposed Methodology Section, we have designed our PathNet ar-

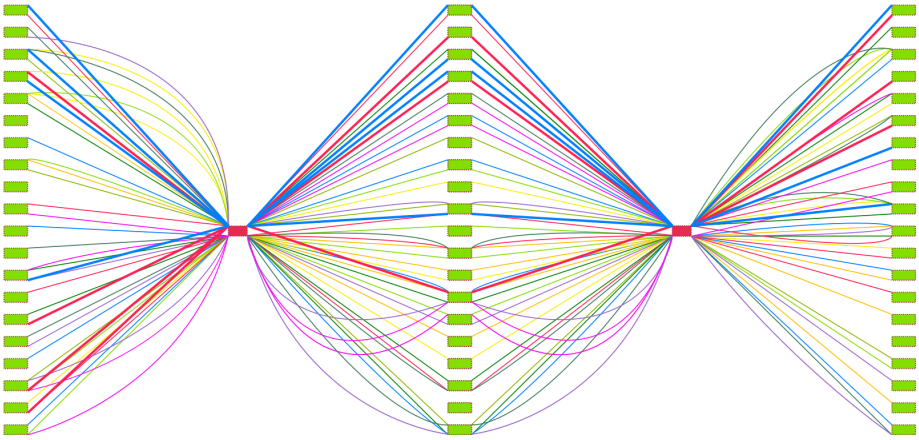

Figure 23: The optimal pathway highlighted by blue colour is achieved when training stage is completed on audio SAVEE (the source data). The paths highlighted by red colour is best pathway gained from the source data.

chitecture including three layers, each layer consists of 20 random modules. However, only up to 4 modules in each layer are activated during training and 20 random pathways are considered (recall that a pathway is a path that connects activated modules in three layers). It may consider all possible configurations that we can come up with those, we used fully-connected layer for each module. Therefore, there is concern regarding the number of parameters in our model since it appears that this architecture may stack up to a large number of parameters, consequently leading to a possibility that the model is prone to over-fitting. However, as explained clearly in Section 3, in each generation only one pathway is trained and another pathway is subsequently trained. This is because we applied a binary tournament selection genetic algorithm. Therefore, the learnable parameters presented in only two pathways are involved in the training stage using stochastic gradient decent for each generation, but this occurs sequentially. Since up to 4 modules are activated in each layer, $64\times64\times3$ is input for the first layer and 20×1 is input for second and third layer, thus the total number of learnable parameters of one module of one pathway is $64\times64\times3\times20 + 20$ for the fist layer, $20\times20 + 20$ for the second and third layer, resulting in up to $(64\times64\times3\times20 + 20 + 20\times20 + 20 + 20\times20 + 20) \times4 = 986{,}480$ parameters for one pathway. While total

number of learnable parameters in AlexNet is 60,954,656. Even if we fine-tune only two last fully-connected layers, then the number of learnable parameters is 16,801,783 which is approximately 17 times greater than the number of learnable parameters for one pathway. In the future work, we may consider the use of convolution operation in each module rather than linear multiplication. This can help capture spatial information and significantly reduce the parameters presented in one pathway.

Additionally, as presented in Section 3.3.2 (Pathway Evolution and Transfer Learning Approach) and also discussed in previous paragraph, in each generation two random pathways are selected among a population of pathways to sequentially train on source dataset. A pathway that has better recognition accuracy is kept and mutated. This means that different emotions (e.g., different labels) are simultaneously used in one pathway and we base on overall performance of one pathway to make a decision whether it is removed or mutated (e.g., competed against a new pathway). If different emotions (e.g., six emotions in this work) are associated with different pathways, for instance one pathway is learned on one emotion, we need a set of six pathways for six emotions and calculate six corresponding recognition accuracy, then we need another set of six pathways. A good point by doing this manner is that when we finish training for a number of generations, we can achieve a best set of six pathways and parameters presented in this set of six pathways are reused and transferred to train simultaneously on six emotions of destination data. Due to six different best sets of parameters from different emotions now involved in one training, we may obtain a better performance overall. However, obviously twelve pathways in each generation are considered, while only two are used in our design of training manner.

Visualization: To understand more how to achieve the best pathways when training our proposed system on the source domain and on the destination domain, we specifically visualize step by step how to obtain such pathways when our model is learned on audio eNTERFACE and the learned emotional knowledge is transferred to the model which is further learned on audio SAVEE as shown in Fig. 19, Fig. 20, Fig. 21, Fig. 22, and Fig. 23.

5. Conclusion

Progress in emotion recognition research has been hindered by the lack of the large amounts of labeled emotion data. To overcome this problem, various studies have widely explored the use of transfer learning approach based on pre-trained/fine-tuning models, however, the proposed approaches have been still suffering from issues such as discarding prior learned information. In this paper, we have proposed utilizing an alternative transfer learning technique using PathNet which is a neural network algorithm that uses agents embedded in the neural network whose task is to discover which parts of the network to reuse for new tasks, leading to successfully addressing the above-mentioned challenges. To verify performance of our proposed architecture, we have conducted various sets of experiments including, transferring the emotional knowledge from one emotion dataset into another, and transferring the learned emotional knowledge from multiple emotion datasets into one another. Experimental results on three datasets: eNTERFACE, SAVEE, and EMO-DB have indicated that our proposed system performs well under the conditions of insufficient emotion data, and significantly better than the recent transfer learning techniques exploiting fine-tuning/pre-trained models.

6. Acknowledgement

This research was supported by an Australian Research Council (ARC) Discovery grant DP140100793.

References

Abbasnejad, I., Sridharan, S., Nguyen, D., Denman, S., Fookes, C., Lucey, S., 2017. Using synthetic data to improve facial expression analysis with 3d convolutional networks, in: 2017 IEEE International Conference on Computer Vision Workshops (ICCVW), pp. 1609–1618.

Badshah, A.M., Ahmad, J., Rahim, N., Baik, S.W., 2017. Speech emotion recognition from spectrograms with deep convolutional neural network, in: 2017 International Conference on Platform Technology and Service (PlatCon), pp. 1–5.

Burkhardt, F., Paeschke, A., Rolfes, M., Sendlmeier, W., Weiss, B., 2005. A database of german emotional speech, in: in Proceedings of Interspeech, Lissabon, pp. 1517–1520.

Chang, J., Scherer, S., 2017. Learning representations of emotional speech with deep convolutional generative adversarial networks, in: 2017 ICASSP, pp. 2746–2750.

Deng, J., Frhholz, S., Zhang, Z., Schuller, B., 2017a. Recognizing emotions from whispered speech based on acoustic feature transfer learning. IEEE Access 5, 5235–5246.

Deng, J., Xu, X., Zhang, Z., Frhholz, S., Schuller, B., 2017b. Universum autoencoder-based domain adaptation for speech emotion recognition. IEEE Signal Processing Letters 24, 500–504.

Eyben, F., 2016. Real-time speech and music classification by large audio feature space extraction. Springer, New York, NY, USA.

Eyben, F., Scherer, K.R., Schuller, B.W., Sundberg, J., Andr, E., Busso, C., Devillers, L.Y., Epps, J., Laukka, P., Narayanan, S.S., Truong, K.P., 2016. The geneva minimalistic acoustic parameter set (gemaps) for voice research and affective computing. IEEE Transactions on Affective Computing 7, 190–202.

Fan, Y., Lu, X., Li, D., Liu, Y., 2016. Video-based emotion recognition using CNN-RNN and C3D hybrid networks, in: Proceedings of the 18th ACM International Conference on Multimodal Interaction, ACM, New York, NY, USA. pp. 445–450.

Fernando, C., Banarse, D., Blundell, C., Zwols, Y., Ha, D., Rusu, A.A., Pritzel, A., Wierstra, D., 2017. Pathnet: Evolution channels gradient descent in super neural networks. CoRR abs/1701.08734.

Gideon, J., Khorram, S., Aldeneh, Z., Dimitriadis, D., Provost, E.M., 2017. Progressive neural networks for transfer learning in emotion recognition. CoRR abs/1706.03256.

Haq, S., Jackson, P., 2010. Machine Audition: Principles, Algorithms and Systems. IGI Global, Hershey PA. chapter Multimodal Emotion Recognition. pp. 398–423.

Harvey, I., 2011. The microbial genetic algorithm, in: 2009 ECRL, Springer-Verlag, Berlin, Heidelberg. pp. 126–133.

Hasani, B., Mahoor, M.H., 2017. Facial expression recognition using enhanced deep 3D convolutional neural networks, in: 2017 CVPRW, pp. 2278–2288.

Hermansky, H., Morgan, N., Bayya, A., Kohn, P., 1992. Rasta-plp speech analysis technique, in: 1992 ICASSP, IEEE Computer Society, Washington, DC, USA. pp. 121–124.

Kaya, H., Grpnar, F., Salah, A.A., 2017. Video-based emotion recognition in the wild using deep transfer learning and score fusion. Image and Vision Computing 65, 66 – 75. Multimodal Sentiment Analysis and Mining in the Wild Image and Vision Computing.

Kim, J., Englebienne, G., Truong, K.P., Evers, V., 2017a. Deep temporal models using identity skip-connections for speech emotion recognition, in: Proceedings of ACM Multimedia, pp. 1006–1013.

Kim, J., Englebienne, G., Truong, K.P., Evers, V., 2017b. Towards speech emotion recognition 'in the wild' using aggregated corpora and deep multi-task learning, in: Proceedings of the INTERSPEECH, pp. 1113–1117.

Kim, J., Truong, K., Englebienne, G., Evers, V., 2017c. Learning spectro-temporal features with 3DCNNs for speech emotion recognition, in: Proceedings of International Conference on Affective Computing and Intelligent Interaction, pp. 383–388.

Krizhevsky, A., Hinton, G., 2009. Learning multiple layers of features from tiny images. Master's thesis, Department of Computer Science, University of Toronto .

Latif, S., Rana, R., Younis, S., Qadir, J., Epps, J., 2018. Cross corpus speech emotion classification- an effective transfer learning technique. CoRR abs/1801.06353.

LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. NATURE 521, 436–444.

Lee, S.W., Kim, J.H., Jun, J., Ha, J.W., Zhang, B.T., 2017. Overcoming catastrophic forgetting by incremental moment matching, in: Guyon, I., Luxburg, U.V., Bengio,

S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (Eds.), 2017 NIPS. Curran Associates, Inc., pp. 4655–4665.

Martin, O., Kotsia, I., Macq, B., Pitas, I., 2006. The eNTERFACE' 05 audio-visual emotion database, in: Data Engineering Workshops, 2006. Proceedings. 22nd International Conference on, pp. 8–8.

Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y., 2011. Reading digits in natural images with unsupervised feature learning, in: NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011, pp. 1–3.

Ng, H.W., Nguyen, V.D., Vonikakis, V., Winkler, S., 2015. Deep learning for emotion recognition on small datasets using transfer learning, in: ICMI, ACM, New York, NY, USA. pp. 443–449.

Nguyen, D., Nguyen, K., Sridharan, S., Abbasnejad, I., Dean, D., Fookes, C., 2018. Meta transfer learning for facial emotion recognition, in: ICPR, pp. 3543–3548.

Nguyen, D., Nguyen, K., Sridharan, S., Dean, D., Fookes, C., 2018. Deep spatio-temporal feature fusion with compact bilinear pooling for multimodal emotion recognition. Computer Vision and Image Understanding .

Nguyen, D., Nguyen, K., Sridharan, S., Ghasemi, A., Dean, D., Fookes, C., 2017. Deep spatio-temporal features for multimodal emotion recognition, in: 2017 IEEE Winter Conference on Applications of Computer Vision (WACV), pp. 1215–1223.

Nguyen, D., Sridharan, S., Nguyen, D.T., Denman, S., Tran, S.N., Zeng, R., Fookes, C., 2020. Joint deep cross-domain transfer learning for emotion recognition. arXiv:2003.11136.

Pao, T.L., Chen, Y.T., Yeh, J.H., Li, P.J., 2006. Mandarin emotional speech recognition based on SVM and NN, in: 18th International Conference on Pattern Recognition (ICPR'06), pp. 1096–1100.

Rabiner, L.R., Sambur, M.R., 1975. An algorithm for determining the endpoints of isolated utterances. The Bell System Technical Journal 54, 297–315.

Rusu, A.A., Rabinowitz, N.C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., Hadsell, R., 2016. Progressive neural networks. CoRR abs/1606.04671.

Sahu, S., Gupta, R., Sivaraman, G., Espy-Wilson, C., AbdAlmageed, W., 2017. Adversarial auto-encoders for speech based emotion recognition, in: InterSpeech 2017, pp. 1243–1247.

Shen, P., Changjun, Z., Chen, X., 2011. Automatic speech emotion recognition using support vector machine, in: Proceedings of 2011 International Conference on Electronic Mechanical Engineering and Information Technology, pp. 621–625.

Sokolova, M., Lapalme, G., 2009. A systematic analysis of performance measures for classification tasks. Information Processing & Management 45, 427 – 437.

Ververidis, D., Kotropoulos, C., Pitas, I., 2004. Automatic emotional speech classification, in: 2004 IEEE International Conference on Acoustics, Speech, and Signal Processing, pp. I–593–6 vol.1.

Viola, P., Jones, M.J., 2004. Robust real-time face detection. Int. J. Comput. Vision 57, 137–154.

Xiao, Z., Dellandrea, E., Dou, W., Chen, L., 2005. Features extraction and selection for emotional speech classification, in: IEEE Conference on Advanced Video and Signal Based Surveillance, 2005., pp. 411–416.

Zhang, S., Huang, T., Gao, W., 2018. Speech emotion recognition using deep convolutional neural network and discriminant temporal pyramid matching. IEEE Transactions on Multimedia 20, 1576–1590.

Zhang, S., Huang, T., Gao, W., Tian, Q., 2017. Learning affective features with a hybrid deep model for audio-visual emotion recognition. IEEE Transactions on Circuits and Systems for Video Technology PP, 1–1.