

Facial Expression Recognition Using Enhanced Deep 3D Convolutional Neural Networks

Behzad Hasani and Mohammad H. Mahoor
Department of Electrical and Computer Engineering
University of Denver, Denver, CO

behzad.hasani@du.edu and mmahoor@du.edu

Abstract

Deep Neural Networks (DNNs) have shown to outperform traditional methods in various visual recognition tasks including Facial Expression Recognition (FER). In spite of efforts made to improve the accuracy of FER systems using DNN, existing methods still are not generalizable enough in practical applications. This paper proposes a 3D Convolutional Neural Network method for FER in videos. This new network architecture consists of 3D Inception-ResNet layers followed by an LSTM unit that together extracts the spatial relations within facial images as well as the temporal relations between different frames in the video. Facial landmark points are also used as inputs to our network which emphasize on the importance of facial components rather than the facial regions that may not contribute significantly to generating facial expressions. Our proposed method is evaluated using four publicly available databases in subject-independent and cross-database tasks and outperforms state-of-the-art methods.

1. Introduction

Facial expressions are one of the most important non-verbal channels for expressing internal emotions and intentions. Ekman *et al.* [13] defined six expressions (viz. anger, disgust, fear, happiness, sadness, and surprise) as basic emotional expressions which are universal among human beings. Automated Facial Expression Recognition (FER) has been a topic of study for decades. Although there have been many breakthroughs in developing automatic FER systems, majority of the existing methods either show undesirable performance in practical applications or lack generalization due to the controlled condition in which they are developed [47].

The FER problem becomes even more difficult when we recognize expressions in videos. Facial expressions have a dynamic pattern that can be divided into three phases: on-

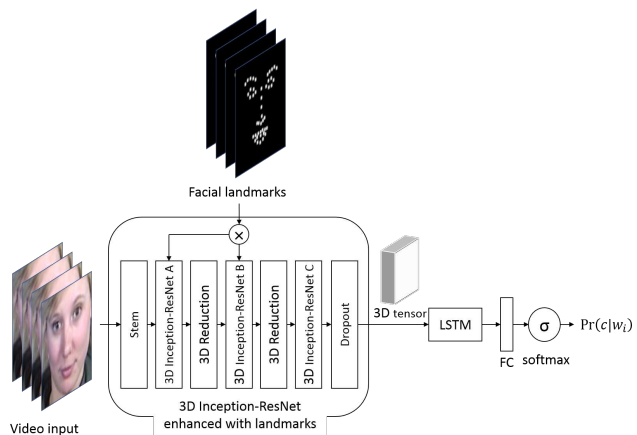


Figure 1. Proposed method

set, peak and offset, where the onset describes the beginning of the expression, the peak (aka apex) describes the maximum intensity of the expression and the offset describes the moment when the expression vanishes. Most of the times, the entire event of facial expression from the onset to the offset is very quick, which makes the process of expression recognition very challenging [55].

Many methods have been proposed for automated facial expression recognition. Most of the traditional approaches mainly consider still images independently while ignore the temporal relations of the consecutive frames in a sequence which are essential for recognizing subtle changes in the appearance of facial images especially in transiting frames between emotions. Recently, with the help of Deep Neural Networks (DNNs), more promising results are reported in the field [38, 39]. While in traditional approaches engineered features are used to train classifiers, DNNs have the ability to extract more discriminative features which yield in a better interpretation of the texture of human face in visual data.

One of the problems in FER is that training neural net-

works is significantly more difficult as most of the existing databases have a small number of images or video sequences for certain emotions [39]. Also, most of these databases contain still images that are unrelated to each other (instead of having consecutive frames of exhibiting the expression from onset to offset) which makes the task of sequential image labeling more difficult.

In this paper, we propose a method which extracts temporal relations of consecutive frames in a video sequence using 3D convolutional networks and Long Short-Term Memory (LSTM). Furthermore, we extract and incorporate facial landmarks in our proposed method that emphasize on more expressive facial components which improve the recognition of subtle changes in the facial expressions in a sequence (Figure 1). We evaluate our proposed method using four well-known facial expression databases (CK+, MMI, FERA, and DISFA) in order to classify the expressions. Furthermore, we examine the ability of our method in recognition of facial expressions in cross-database classification tasks.

The remainder of the paper is organized as follows: Section 2 provides an overview of the related work in this field. Section 3 explains the network proposed in this research. Experimental results and their analysis are presented in Section 4 and finally the paper is concluded in Section 5.

2. Related work

Traditionally, algorithms for automated facial expression recognition consist of three main modules, viz. registration, feature extraction, and classification. Detailed survey of different approaches in each of these steps can be found in [44]. Conventional algorithms for affective computing from faces use engineered features such as Local Binary Patterns (LBP) [47], Histogram of Oriented Gradients (HOG) [8], Local Phase Quantization (LPQ) [59], Histogram of Optical Flow [9], facial landmarks [6, 7], and PCA-based methods [36]. Since the majority of these features are hand-crafted for their specific application of recognition, they often lack required generalizability in cases where there is high variation in lighting, views, resolution, subjects’ ethnicity, etc.

One of the effective approaches for achieving better recognition rates for sequence labeling task is to extract the temporal relations of frames in a sequence. Extracting these temporal relations has been studied using traditional methods in the past. Examples of these attempts are Hidden Markov Models [5, 60, 64] (which combine temporal information and apply segmentation on videos), Spatio-Temporal Hidden Markov Models (ST-HMM) by coupling S-HMM and T-HMM [50], Dynamic Bayesian Networks (DBN) [45, 63] associated with a multi-sensory information fusion strategy, Bayesian temporal models [46] to capture the dynamic facial expression transition, and Conditional Random Fields (CRFs) [19, 20, 25, 48] and their

extensions such as Latent-Dynamic Conditional Random Fields (LD-CRFs) and Hidden Conditional Random Fields (HCRFs) [58].

In recent years, “Convolutional Neural Networks” (CNNs) have become the most popular approach among researchers in the field. AlexNet [27] is based on the traditional CNN layered architecture which consists of several convolution layers followed by max-pooling layers and Rectified Linear Units (ReLUs). Szegedy *et al.* [52] introduced GoogLeNet which is composed of multiple “Inception” layers. Inception applies several convolutions on the feature map in different scales. Mollahosseini *et al.* [38, 39] have used the Inception layer for the task of facial expression recognition and achieved state-of-the-art results. Following the success of Inception layers, several variations of them have been proposed [24, 53]. Moreover, Inception layer is combined with residual unit introduced by He *et al.* [21] and it shows that the resulting architecture accelerates the training of Inception networks significantly [51].

One of the major restrictions of ordinary Convolutional Neural Networks is that they only extract spatial relations of the input data while ignore the temporal relations of them if they are part of a sequenced data. To overcome this problem, 3D Convolutional Neural Networks (3D-CNNs) have been proposed. 3D-CNNs slide over the temporal dimension of the input data as well as the spatial dimension enabling the network to extract feature maps containing temporal information which is essential for sequence labeling tasks. Song *et al.* [49] have used 3D-CNNs for 3D object detection task. Molchanov *et al.* [37] have proposed a recurrent 3D-CNN for dynamic hand gesture recognition and Fan *et al.* [15] won the EmotiW 2016 challenge by cascading 3D-CNNs with LSTMs.

Traditional Recurrent Neural Networks (RNNs) can learn temporal dynamics by mapping input sequences to a sequence of hidden states, and also mapping the hidden states to outputs [12]. Although RNNs have shown promising performance on various tasks, it is not easy for them to learn long-term sequences. This is mainly due to the vanishing/exploding gradients problem [23] which can be solved by having a memory for remembering and forgetting the previous states. LSTMs [23] provide such memory and can memorize the context information for long periods of time. LSTM modules have three gates: 1) the input gate (*i*) 2) the forget gate (*f*) and 3) the output gate (*o*) which overwrite, keep, or retrieve the memory cell *c* respectively at the timestep *t*. Letting $\sigma(x) = (1 + \exp(-x))^{-1}$ be the *sigmoid* function and $\phi(x) = \frac{\exp(x) - \exp(-x)}{\exp(x) + \exp(-x)} = 2\sigma(2x) - 1$ be the *hyperbolic tangent* function. Letting *x*, *h*, *c*, *W*, and *b* be the input, output, cell state, parameter matrix, and parameter vector respectively. The LSTM updates for the timestep *t* given inputs *x_t*, *h_{t-1}*, and *c_{t-1}* are as follows:

$$\begin{aligned}
f_t &= \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \\
i_t &= \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \\
o_t &= \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \\
g_t &= \phi(W_C \cdot [h_{t-1}, x_t] + b_C) \\
C_t &= f_t * C_{t-1} + i_t * g_t \\
h_t &= o_t * \phi(C_t)
\end{aligned} \tag{1}$$

Several works have used LSTMs for the task of sequence labeling. Byeon *et al.* [3] proposed an LSTM-based network applying LSTMs in four direction sliding windows and achieved impressive results. Fan *et al.* [15] cascaded 2D-CNN with LSTMs and combined the feature map with 3D-CNNs for facial expression recognition. Donahue *et al.* [12] proposed Long-term Recurrent Convolutional Network (LRCN) by combining CNNs and LSTMs which is both spatially and temporally deep and has the flexibility to be applied to different vision tasks involving sequential inputs and outputs.

3. Proposed method

While Inception and ResNet have shown remarkable results in FER [20, 52], these methods do not extract the temporal relations of the input data. Therefore, we propose a 3D Inception-ResNet architecture to address this issue. Our proposed method, extracts both spatial and temporal features of the sequences in an end-to-end neural network. Another component of our method is incorporating facial landmarks in an automated manner during training in the proposed neural network. These facial landmarks help the network to pay more attention to the important facial components in the feature maps which results in a more accurate recognition. The final part of our proposed method is an LSTM unit which takes the enhanced feature map resulted from the 3D Inception-ResNet (3DIR) layer as an input and extracts the temporal information from it. The LSTM unit is followed by a fully-connected layer associated with a softmax activation function. In the following, we explain each of the aforementioned units in detail.

3.1. 3D Inception-ResNet (3DIR)

We propose 3D version of Inception-ResNet network which is slightly shallower than the original Inception-ResNet network proposed in [51]. This network is the result of investigating several variations of Inception-ResNet module and achieves better recognition rates comparing to our other attempts in several databases.

Figure 2 shows the structure of our 3D Inception-ResNet network. The input videos with the size $10 \times 299 \times 299 \times 3$ (10 frames, 299×299 frame size and 3 color channels) are followed by the “stem” layer. Afterwards, stem is followed

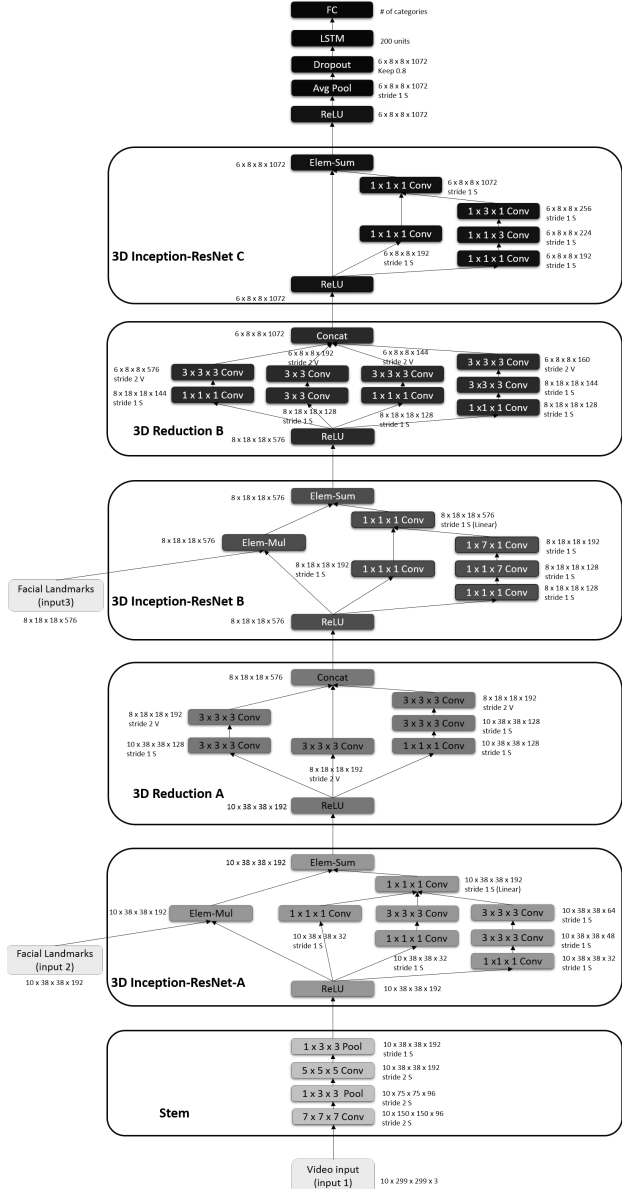


Figure 2. Network architecture. The “V” and “S” marked layers represent “Valid” and “Same” paddings respectively. The size of the output tensor is provided next to each layer.

by 3DIR-A, Reduction-A (which reduces the grid size from 38×38 to 18×18), 3DIR-B, Reduction-B (which reduces the grid size from 18×18 to 8×8), 3DIR-C, Average Pooling, Dropout, and a fully-connected layer respectively. In Figure 2, detailed specification of each layer is provided. Various filter sizes, paddings, strides, and activations have been investigated and the one that had the best performance is presented in this paper.

We should mention that all convolution layers (except the ones that are indicated as “Linear” in Figure 2) are fol-

lowed by an ReLU [27] activation function to avoid the vanishing gradient problem.

3.2. Facial landmarks

As mentioned before, the main reason we use facial landmarks in our network is to differentiate between the importance of main facial components (such as eyebrows, lip corners, eyes, etc.) and other parts of the face which are less expressive of facial expressions. As oppose to general object recognition task, in FER, we have the advantage of extracting facial landmarks and using this information to improve the recognition rate. In a similar approach, Jaiswal *et al.* [26] proposed incorporation of binary masks around different parts of the face in order to encode the shape of different face components. However, in this work authors perform AU recognition by using CNN as a feature extractor for training Bi-directional Long Short-Term Memory while in our approach, we preserve the temporal order of the frames throughout the network and train CNN and LSTMs simultaneously in an end-to-end network. We incorporate the facial landmarks by replacing the shortcut in residual unit on original ResNet with element-wise multiplication of facial landmarks and the input tensor of the residual unit (Figures 1 and 2).

In order to extract the facial landmarks, OpenCV face recognition is used to obtain bounding boxes of the faces. A face alignment algorithm via regression local binary features [41, 61] was used to extract 66 facial landmark points. The facial landmark localization technique was trained using the annotations provided from the 300W competition [42, 43].

After detecting and saving the facial landmarks for all of the databases, the facial landmark filters are generated for each sequence automatically during training phase. Given the facial landmarks for each frame of a sequence, we initially resize all of the images in the sequence to their corresponding filter size in the network. Afterwards, we assign weights to all of the pixels in a frame of a sequence based on their distances to the detected landmarks. The closer a pixel is to a facial landmark, the greater weight is assigned to that pixel. After investigating several distance measures, we concluded that *Manhattan* distance with a linear weight function results in a better recognition rate in various databases. The Manhattan distance between two items is the sum of the differences of their corresponding components (in this case two components).

The weight function that we defined to assign the weight values to their corresponding feature is a simple linear function of the Manhattan distance defined as follows:

$$\omega(L, P) = 1 - 0.1 \cdot d_{M(L,P)} \quad (2)$$

where $d_{M(L,P)}$ is the Manhattan distance between the facial landmark L and pixel P . Therefore, places in which facial




Figure 3. Sample image from MMI database (left) and its corresponding filter in the network (right). Best in color.

landmarks are located will have the highest value and their surrounding pixels will have lower weights proportional to their distance from the corresponding facial landmark. In order to avoid overlapping between two adjacent facial landmarks, we define a 7×7 window around each facial landmark and apply the weight function for these 49 pixels for each landmark separately. Figure 3 shows an example of facial image from MMI database and its corresponding facial landmark filter in the network. We do not incorporate the facial landmarks with the third 3D Inception-ResNet module since the resulting feature map size at this stage becomes very small for calculating facial landmark filter.

Incorporating facial landmarks in our network replaces the shortcut in original ResNets [22] with the element-wise multiplication of the weight function ω and input layer x_l as follows:

$$\begin{aligned} y_l &= \omega(L, P) \circ x_l + F(x_l, W_l) \\ x_{l+1} &= f(y_l) \end{aligned} \quad (3)$$

where x_l and x_{l+1} are input and output of the l -th layer, \circ is Hadamard product symbol, F is a residual function (in our case Inception layer convolutions), and f is an activation function.

3.3. Long Short-Term Memory unit

As explained earlier, to capture the temporal relations of the resulted feature map from 3DIR and take these relations into account by the time of classifying the sequences in the softmax layer, we used an LSTM unit as it is shown in Figure 2. Using the LSTM unit makes perfect sense since the resulted feature map from the 3DIR unit contains the time notion of the sequences within the feature map. Therefore, vectorizing the resulting feature map of 3DIR on its sequence dimension, will provide the required sequenced input for the LSTM unit. While other still image LSTM-based methods, a vectorized non-sequenced feature map (which obviously does not contain any time notion) is fed to the LSTM unit, our method saves the time order of the input sequences and passes this feature map to the LSTM unit. We investigated that 200 hidden units for the LSTM unit is a reasonable amount for the task of FER (Figure 2).

The proposed network was implemented using a combination of TensorFlow [1] and TFlern [10] toolboxes on NVIDIA Tesla K40 GPUs. In the training phase we used asynchronous stochastic gradient descent with momentum of 0.9, weight decay of 0.0001, and learning rate of 0.01. We used categorical cross entropy as our loss function and accuracy as our evaluation metric.

4. Experiments and results

In this section, we briefly review the databases we used for evaluating our method. We then report the results of our experiments using these databases and compare the results with the state of the arts.

4.1. Face databases

Since our method is designed mainly for classifying sequences of inputs, databases that contain only independent unrelated still images of facial expressions such as MultiPie [18], SFEW [11], FER2013 [17] cannot be examined by our method. We evaluate our proposed method on MMI [40], extended CK+ [32], GEMEP-FERA [2], and DISFA [33] which contain videos of annotated facial expressions. In the following, we briefly review the contents of these databases.

MMI: The MMI [40] database contains more than 20 subjects, ranging in age from 19 to 62, with different ethnicities (European, Asian, or South American). In MMI, the subjects' facial expressions start from the neutral state to the apex of one of the six basic facial expressions and then returns to the neutral state again. Subjects were instructed to display 79 series of facial expressions, six of which are prototypic emotions (angry, disgust, fear, happy, sad, and surprise). We extracted static frames from each sequence, which resulted in 11,500 images. Afterwards, we divided videos into sequences of ten frames to shape the input tensor for our network.

CK+: The extended Cohn-Kanade database (CK+) [32] contains 593 videos from 123 subjects. However, only 327 sequences from 118 subjects contain facial expression labels. Sequences in this database start from the neutral state and end at the apex of one of the six basic expressions (angry, contempt, disgust, fear, happy, sad, and surprise). CK+ primarily contains frontal face poses only. In order to make the database compatible with our network, we consider the last ten frames of each sequence as an input sequence in our network.

FERA: The GEMEP-FERA database [2] is a subset of the GEMEP corpus used as database for the FERA 2011 challenge [56] developed by the Geneva Emotion Research Group at the University of Geneva. This database contains 87 image sequences of 7 subjects. Each subject shows facial expressions of the emotion categories: Anger, Fear, Joy,

Relief, and Sadness. Head pose is primarily frontal with relatively fast movements. Each video is annotated with AUs and holistic expressions. By extracting static frames from the sequences, we obtained around 7,000 images. We divided these emotion videos into sequences of ten frames to shape the input tensor for our network.

DISFA: Denver Intensity of Spontaneous Facial Actions (DISFA) database [33] is one of a few naturalistic databases that have been FACS coded by AU intensity values. This database consists of 27 subjects. The subjects are asked to watch YouTube videos while their spontaneous facial expressions are recorded. Twelve AUs are coded for each frame and AU intensities are on a six-point scale between 0-5, where 0 denotes the absence of the AU, and 5 represents maximum intensity. As DISFA is not emotion-specified coded, we used EMFACS system [16] to convert AU FACS codes to seven expressions (angry, disgust, fear, happy, neutral, sad, and surprise) which resulted in around 89,000 images in which the majority have neutral expressions. Same as other databases, we divided the videos of emotions into sequences of ten frames to shape the input tensor for our network.

4.2. Results

As mentioned earlier, after detecting faces we extract 66 facial landmark points by a face alignment algorithm via regression local binary features. Afterwards, we resize the faces to 299×299 pixels. One of the reasons why we choose large image size as input is the fact that larger images and sequences will enable us to have deeper networks and extract more abstract features from sequences. All of the networks have the same settings (shown in Figure 2 in detail) and are trained from scratch for each database separately.

We evaluate the accuracy of our proposed method with two different sets of experiments: "subject-independent" and "cross-database" evaluations.

4.2.1 Subject-independent task

In the subject-independent task, each database is split into training and validation sets in a strict subject independent manner. In all databases, we report the results using the 5-fold cross-validation technique and then averaging the recognition rates over five folds. For each database and each fold, we trained our proposed network entirely from scratch with the aforementioned settings. Table 1 shows the recognition rates achieved on each database in the subject-independent case and compares the results with the state-of-the-art methods. In order to compare the impact of incorporating facial landmarks, we also provide the results of our network while the landmark multiplication unit is removed and replaced with a simple shortcut between the input and output of the residual unit. In this case, we randomly se-

lect 20 percent of the subjects as the test set and report the results on those subjects. Table 1 also provides the recognition rates of the traditional 2D Inception-ResNet from [20] which does not contain facial landmarks and the LSTM unit (DISFA is not experimented in this study).

Comparing the recognition rates of the 3D and 2D Inception-ResNets in Table 1, shows that the sequential processing of facial expressions considerably enhances the recognition rate. This improvement is more apparent in MMI and FERA databases. Incorporating landmarks in the network is proposed to emphasize on more important facial changes over time. Since changes in the lips or eyes are much more expressive than the changes in other components such as the cheeks, we utilize facial landmarks to enhance these temporal changes in the network flow.

The “3D Inception-ResNet with landmarks” column in Table 1 shows the impact of this enhancement in different databases. It can be seen that compared with other networks, there is a considerable improvement in recognition rates especially in FERA and MMI databases. The results on DISFA, however, show higher fluctuations over different folds which can be in part due to the abundance of inactive frames in this database which causes confusion in recognizing different expressions. Therefore, the folds that contain more neutral faces, would show lower recognition rates.

Comparing to other state-of-the-art works, our method outperforms others in FERA and DISFA databases while achieves comparable results in CK+ and MMI databases (Table 1). Most of these works use traditional approaches including hand-crafted features tuned for that specific database, while our network’s settings are the same for all databases. Also, due to the limited number of samples in these databases, it is difficult to properly train a deep neural network and avoid the overfitting problem. For these reasons and in order to have a better understanding about our proposed method, we also experimented the cross-database task.

Figure 4 shows the resulting confusion matrices of our 3D Inception-ResNet with incorporating landmarks on different databases over the 5 folds. On CK+ (Figure 4a), it can be seen that very high recognition rates have been achieved. The recognition rates of happiness, sadness, and surprise are higher than those of other expressions. The highest confusion occurred between the happiness and contempt expressions which can be caused from the low number of contempt sequences in this database (only 18 sequences). On MMI (Figure 4b), a perfect recognition is achieved for the happy expression. It can be seen that there is a high confusion between the sad and fear expressions as well as the angry and sad expressions. Considering the fact that MMI is a highly imbalanced dataset, these confusions are reasonable. On FERA (Figure 4c), the highest and the lowest recognition rates belong to joy and relief respectively. The relief

category in this database has some similarities with other categories especially with joy. These similarities make the classification so difficult even for humans. Despite these challenges, our method has performed well on all of the categories and outperforms state of the arts. On DISFA (Figure 4d), we can see the highest confusion rate compared with other databases. As mentioned earlier, this database contains long inactive frames, which means that the number of neutral sequences is considerably higher than other categories. This imbalanced training data has made the network to be biased toward the neutral category and therefore we can observe a high confusion rate between the neutral expression and other categories in this database. Despite the low number of angry and sad sequences in this database, our method has been able achieve satisfying recognition rates in these categories.

4.2.2 Cross-database task

In the cross-database task, for testing each database, that database is entirely used for testing the network and the rest of the databases are used to train the network. The same network architecture as subject-independent task (Figure 2) was used for this task. Table 2 shows the recognition rate achieved on each database in the cross-database case and it also compares the results with other state-of-the-art methods. It can be seen that our method outperforms the state-of-the-art results in CK+, FERA, and DISFA databases. On MMI, our method does not show improvements comparing to others (e.g. [62]). However, authors in [62] trained their classifier only with CK+ database while our method uses instances from two additional databases (DISFA and FERA) with completely different settings and subjects which add significant amount of ambiguity in the training phase.

In order to have a fair comparison with other methods, we provide the different settings used by the works mentioned in Table 2. The results provided in [34] are achieved by training the models on one of the CK+, MMI, and FEED-TUM databases and tested on the rest. The reported result in [47] is the best achieved results using different SVM kernels trained on CK+ and tested on MMI database. In [35] several experiments were performed using four classifiers (SVM, Nearest Mean Classifier, Weighted Template Matching, and K-nearest neighbors). The reported results in this work for CK+ is trained on MMI and Jaffe databases while the reported results for MMI is trained on the CK+ database only. As mentioned earlier, in [62] a Multiple Kernel Learning algorithm is used and the cross-database experiments are trained on CK+, evaluated on MMI and vice versa. In [38] a DNN network is proposed using traditional Inception layer. The networks for the cross-database case in this work are tested on either CK+, MultiPIE, MMI, DISFA, FERA, SFEW, or FER2013 while trained on the rest. Some

	state-of-the-art methods	2D Inception-ResNet	3D Inception-ResNet	3D Inception-ResNet + landmarks
CK+	84.1 [34], 84.4 [28], 88.5 [54], 92.0 [29], 93.2 [38], 92.4 [30], 93.6 [62]	85.77	89.50	93.21±2.32
MMI	63.4 [30], 75.12 [31], 74.7 [29], 79.8 [54], 86.7 [47], 78.51 [36]	55.83	67.50	77.50±1.76
FERA	56.1 [30], 55.6 [57], 76.7 [38]	49.64	67.74	77.42±3.67
DISFA	55.0 [38]	-	51.35	58.00±5.77

Table 1. Recognition rates (%) in subject-independent task

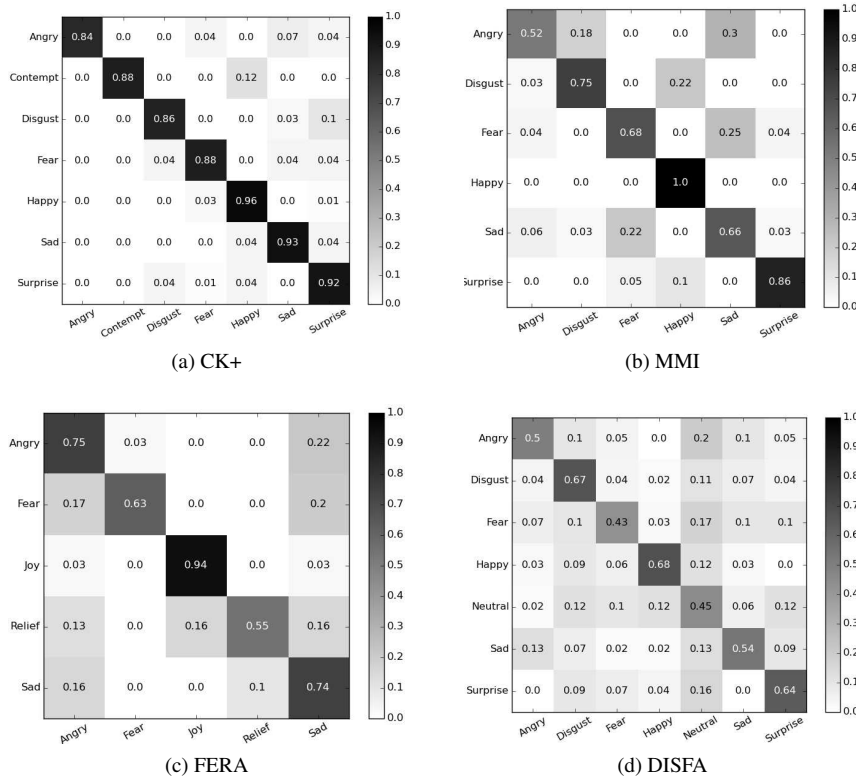


Figure 4. Confusion matrices of 3D Inception-ResNet with landmarks for subject-independent task

of the expressions of these databases are excluded in this study (such as neutral, relief, and contempt). There are other works that perform their experiments on action unit recognition task [4, 14, 26] but since fair comparison of action unit recognition and facial expression recognition is not easily obtainable, we did not mention these works in Tables 1 and 2.

Figure 5 shows the resulting confusion matrices of our experiments on 3D Inception-ResNet with landmarks in cross-database task. For CK+ (Figure 5a), we exclude the contempt sequences in the test phase since other databases that are used for training the network, do not contain contempt category. Except for the fear expression (which has

very few number of samples in other databases), the network has been able to correctly recognize other expressions. For MMI (Figure 5b), highest recognition rate belongs to surprise while the lowest one belongs to fear. Also, we can see high confusion rate in recognizing sadness. On FERA (Figure 5c), we exclude relief category as other databases do not contain this emotion. Considering the fact that only half of the train categories exist in the test set, the network shows acceptable performance in correctly recognizing emotions. However, surprise category has made significant confusion in all of the categories. On DISFA (Figure 5d), we exclude the neutral category as other databases do not contain this category. Highest recognition rates belong to happy and sur-

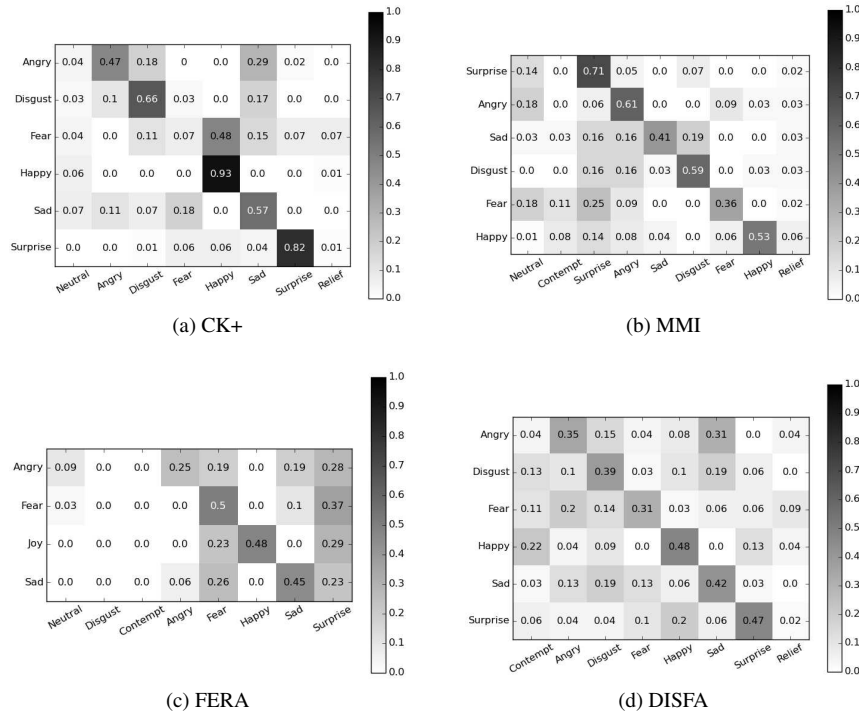


Figure 5. Confusion matrices of 3D Inception-ResNet with landmarks for cross-database task

prise emotions while lowest one belongs to fear. Comparing to other databases, we can see a significant increase in confusion rate in all of the categories. This can be in part due to the fact that emotions in DISFA are “spontaneous” while emotions in the training databases are “posed”. Based on the aforementioned results, our method provides a comprehensive solution that can generalize well to practical applications.

5. Conclusion

In this paper, we presented a 3D Deep Neural Network for the task of facial expression recognition in videos. We

proposed the 3D Inception-ResNet (3DIR) network which extends the well-known 2D Inception-ResNet module for processing image sequences. This additional dimension will result in a volume of feature maps and will extract the spatial relations between frames in a sequence. This module is followed by an LSTM which takes these temporal relations into account and uses this information to classify the sequences. In order to differentiate between facial components and other parts of the face, we incorporated facial landmarks in our proposed method. These landmarks are multiplied with the input tensor in the residual module which is replaced with the shortcuts in the traditional residual layer.

We evaluated our proposed method in subject-independent and cross-database tasks. Four well-known databases were used to evaluate the method: CK+, MMI, FERA, and DISFA. Our experiments show that the proposed method outperforms many of the state-of-the-art methods in both tasks and provides a general solution for the task of FER.

6. Acknowledgement

This work is partially supported by the NSF grants IIS-1111568 and CNS-1427872. We gratefully acknowledge the support from NVIDIA Corporation with the donation of the Tesla K40 GPUs used for this research.

	state-of-the-art methods	3D Inception-ResNet + landmarks
CK+	47.1 [34], 56.0 [35], 61.2 [62], 64.2 [38]	67.52
MMI	51.4 [34], 50.8 [47], 36.8 [35], 55.6 [38], 66.9 [62]	54.76
FERA	39.4[38]	41.93
DISFA	37.7 [38]	40.51

Table 2. Recognition rates (%) in cross-database task

References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016. **5**
- [2] T. Bänziger and K. R. Scherer. Introducing the geneva multimodal emotion portrayal (gemep) corpus. *Blueprint for affective computing: A sourcebook*, pages 271–294, 2010. **5**
- [3] W. Byeon, T. M. Breuel, F. Raue, and M. Liwicki. Scene labeling with lstm recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3547–3555, 2015. **3**
- [4] W.-S. Chu, F. De la Torre, and J. F. Cohn. Modeling spatial and temporal cues for multi-label facial action unit detection. *arXiv preprint arXiv:1608.00911*, 2016. **7**
- [5] I. Cohen, N. Sebe, A. Garg, L. S. Chen, and T. S. Huang. Facial expression recognition from video sequences: temporal and static modeling. *Computer Vision and image understanding*, 91(1):160–187, 2003. **2**
- [6] T. F. Cootes, G. J. Edwards, C. J. Taylor, et al. Active appearance models. *IEEE Transactions on pattern analysis and machine intelligence*, 23(6):681–685, 2001. **2**
- [7] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models-their training and application. *Computer vision and image understanding*, 61(1):38–59, 1995. **2**
- [8] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 886–893. IEEE, 2005. **2**
- [9] N. Dalal, B. Triggs, and C. Schmid. Human detection using oriented histograms of flow and appearance. In *European conference on computer vision*, pages 428–441. Springer, 2006. **2**
- [10] A. Damien et al. Tflern. <https://github.com/tflern/tflern>, 2016. **5**
- [11] A. Dhall, R. Goecke, S. Lucey, and T. Gedeon. Static facial expression analysis in tough conditions: Data, evaluation protocol and benchmark. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*, pages 2106–2112. IEEE, 2011. **5**
- [12] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015. **2, 3**
- [13] P. Ekman and W. V. Friesen. Constants across cultures in the face and emotion. *Journal of personality and social psychology*, 17(2):124, 1971. **1**
- [14] C. Fabian Benitez-Quiroz, R. Srinivasan, and A. M. Martinez. Emotionet: An accurate, real-time algorithm for the automatic annotation of a million facial expressions in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5562–5570, 2016. **7**
- [15] Y. Fan, X. Lu, D. Li, and Y. Liu. Video-based emotion recognition using cnn-rnn and c3d hybrid networks. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction, ICMI 2016*, pages 445–450, New York, NY, USA, 2016. ACM. **2, 3**
- [16] W. V. Friesen and P. Ekman. Emfacs-7: Emotional facial action coding system. *Unpublished manuscript, University of California at San Francisco*, 2(36):1, 1983. **5**
- [17] I. J. Goodfellow, D. Erhan, P. L. Carrier, A. Courville, M. Mirza, B. Hamner, W. Cukierski, Y. Tang, D. Thaler, D.-H. Lee, et al. Challenges in representation learning: A report on three machine learning contests. In *International Conference on Neural Information Processing*, pages 117–124. Springer, 2013. **5**
- [18] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. Multi-pie. *Image and Vision Computing*, 28(5):807–813, 2010. **5**
- [19] B. Hasani, M. M. Arzani, M. Fathy, and K. Raahemifar. Facial expression recognition with discriminatory graphical models. In *2016 2nd International Conference of Signal Processing and Intelligent Systems (ICSPIS)*, pages 1–7, Dec 2016. **2**
- [20] B. Hasani and M. H. Mahoor. Spatio-temporal facial expression recognition using convolutional neural networks and conditional random fields. *arXiv preprint arXiv:1703.06995*, 2017. **2, 3, 6**
- [21] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. **2**
- [22] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. pages 630–645, 2016. **4**
- [23] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. **2**
- [24] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. **2**
- [25] S. Jain, C. Hu, and J. K. Aggarwal. Facial expression recognition with temporal modeling of shapes. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*, pages 1642–1649. IEEE, 2011. **2**
- [26] S. Jaiswal and M. Valstar. Deep learning the dynamic appearance and shape of facial action units. In *Applications of Computer Vision (WACV), 2016 IEEE Winter Conference on*, pages 1–8. IEEE, 2016. **4, 7**
- [27] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. **2, 4**
- [28] S. H. Lee, K. N. K. Plataniotis, and Y. M. Ro. Intra-class variation reduction using training expression images for sparse representation based facial expression recognition. *IEEE Transactions on Affective Computing*, 5(3):340–351, 2014. **7**
- [29] M. Liu, S. Li, S. Shan, and X. Chen. Au-aware deep networks for facial expression recognition. In *Automatic Face and Gesture Recognition (FG), 2013 10th IEEE International Conference and Workshops on*, pages 1–6. IEEE, 2013. **7**

- [30] M. Liu, S. Li, S. Shan, R. Wang, and X. Chen. Deeply learning deformable facial action parts model for dynamic expression analysis. In *Asian Conference on Computer Vision*, pages 143–157. Springer, 2014. 7
- [31] M. Liu, S. Shan, R. Wang, and X. Chen. Learning expressionlets on spatio-temporal manifold for dynamic facial expression recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1749–1756, 2014. 7
- [32] P. Lucey, J. F. Cohn, T. Kanade, J. Saragih, Z. Ambadar, and I. Matthews. The extended cohn-kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2010 IEEE Computer Society Conference on*, pages 94–101. IEEE, 2010. 5
- [33] S. M. Mavadati, M. H. Mahoor, K. Bartlett, P. Trinh, and J. F. Cohn. Disfa: A spontaneous facial action intensity database. *IEEE Transactions on Affective Computing*, 4(2):151–160, 2013. 5
- [34] C. Mayer, M. Eggers, and B. Radig. Cross-database evaluation for facial expression recognition. *Pattern recognition and image analysis*, 24(1):124–132, 2014. 6, 7, 8
- [35] Y.-Q. Miao, R. Araujo, and M. S. Kamel. Cross-domain facial expression recognition using supervised kernel mean matching. In *Machine Learning and Applications (ICMLA), 2012 11th International Conference on*, volume 2, pages 326–332. IEEE, 2012. 6, 8
- [36] M. Mohammadi, E. Fatemizadeh, and M. H. Mahoor. Pca-based dictionary building for accurate facial expression recognition via sparse representation. *Journal of Visual Communication and Image Representation*, 25(5):1082–1092, 2014. 2, 7
- [37] P. Molchanov, X. Yang, S. Gupta, K. Kim, S. Tyree, and J. Kautz. Online detection and classification of dynamic hand gestures with recurrent 3d convolutional neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4207–4215, 2016. 2
- [38] A. Mollahosseini, D. Chan, and M. H. Mahoor. Going deeper in facial expression recognition using deep neural networks. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–10. IEEE, 2016. 1, 2, 6, 7, 8
- [39] A. Mollahosseini, B. Hasani, M. J. Salvador, H. Abdollahi, D. Chan, and M. H. Mahoor. Facial expression recognition from world wild web. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2016. 1, 2
- [40] M. Pantic, M. Valstar, R. Rademaker, and L. Maat. Web-based database for facial expression analysis. In *2005 IEEE international conference on multimedia and Expo*, pages 5–pp. IEEE, 2005. 5
- [41] S. Ren, X. Cao, Y. Wei, and J. Sun. Face alignment at 3000 fps via regressing local binary features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1685–1692, 2014. 4
- [42] C. Sagonas, E. Antonakos, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic. 300 faces in-the-wild challenge: Database and results. *Image and Vision Computing*, 47:3–18, 2016. 4
- [43] C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic. A semi-automatic methodology for facial landmark annotation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 896–903, 2013. 4
- [44] E. Sariyanidi, H. Gunes, and A. Cavallaro. Automatic analysis of facial affect: A survey of registration, representation, and recognition. *IEEE transactions on pattern analysis and machine intelligence*, 37(6):1113–1133, 2015. 2
- [45] N. Sebe, M. S. Lew, Y. Sun, I. Cohen, T. Gevers, and T. S. Huang. Authentic facial expression analysis. *Image and Vision Computing*, 25(12):1856–1863, 2007. 2
- [46] C. Shan, S. Gong, and P. W. McOwan. Dynamic facial expression recognition using a bayesian temporal manifold model. In *BMVC*, pages 297–306. Citeseer, 2006. 2
- [47] C. Shan, S. Gong, and P. W. McOwan. Facial expression recognition based on local binary patterns: A comprehensive study. *Image and Vision Computing*, 27(6):803–816, 2009. 1, 2, 6, 7, 8
- [48] C. Sminchisescu, A. Kanaujia, and D. Metaxas. Conditional models for contextual human motion recognition. *Computer Vision and Image Understanding*, 104(2):210–220, 2006. 2
- [49] S. Song and J. Xiao. Deep sliding shapes for amodal 3d object detection in rgb-d images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 808–816, 2016. 2
- [50] Y. Sun, X. Chen, M. Rosato, and L. Yin. Tracking vertex flow and model adaptation for three-dimensional spatiotemporal face analysis. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 40(3):461–474, 2010. 2
- [51] C. Szegedy, S. Ioffe, and V. Vanhoucke. Inception-v4, inception-resnet and the impact of residual connections on learning. *arXiv preprint arXiv:1602.07261*, 2016. 2, 3
- [52] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015. 2, 3
- [53] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [54] S. Taheri, Q. Qiu, and R. Chellappa. Structure-preserving sparse decomposition for facial expression analysis. *IEEE Transactions on Image Processing*, 23(8):3590–3603, 2014. 7
- [55] Y. Tian, T. Kanade, and J. F. Cohn. Recognizing lower face action units for facial expression analysis. In *Automatic Face and Gesture Recognition, 2000. Proceedings. Fourth IEEE International Conference on*, pages 484–490. IEEE, 2000. 1
- [56] M. F. Valstar, B. Jiang, M. Mehu, M. Pantic, and K. Scherer. The first facial expression recognition and analysis challenge. In *Automatic Face & Gesture Recognition and Workshops (FG 2011), 2011 IEEE International Conference on*, pages 921–926. IEEE, 2011. 5

- [57] M. F. Valstar, B. Jiang, M. Mehu, M. Pantic, and K. Scherer. The first facial expression recognition and analysis challenge. In *Automatic Face & Gesture Recognition and Workshops (FG 2011), 2011 IEEE International Conference on*, pages 921–926. IEEE, 2011. 7
- [58] S. B. Wang, A. Quattoni, L.-P. Morency, D. Demirdjian, and T. Darrell. Hidden conditional random fields for gesture recognition. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 2, pages 1521–1527. IEEE, 2006. 2
- [59] Z. Wang and Z. Ying. Facial expression recognition based on local phase quantization and sparse representation. In *Natural Computation (ICNC), 2012 Eighth International Conference on*, pages 222–225. IEEE, 2012. 2
- [60] M. Yeasin, B. Bulot, and R. Sharma. Recognition of facial expressions and measurement of levels of interest from video. *Multimedia, IEEE Transactions on*, 8(3):500–508, 2006. 2
- [61] L. Yu. face-alignment-in-3000fps. <https://github.com/yulequan/face-alignment-in-3000fps>, 2016. 4
- [62] X. Zhang, M. H. Mahoor, and S. M. Mavadati. Facial expression recognition using $\{1\}$ - $\{p\}$ -norm mkl multiclass-svm. *Machine Vision and Applications*, 26(4):467–483, 2015. 6, 7, 8
- [63] Y. Zhang and Q. Ji. Active and dynamic information fusion for facial expression understanding from image sequences. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 27(5):699–714, 2005. 2
- [64] Y. Zhu, L. C. De Silva, and C. C. Ko. Using moment invariants and hmm in facial expression recognition. *Pattern Recognition Letters*, 23(1):83–91, 2002. 2