EMOTION ESTIMATION FROM VIDEO FOOTAGE WITH LSTM

Samer Attrah* Engineering and Automotive Academy Hogeschool Van Arnhem en Nijmegen samiratra95@gmail.com

April 22, 2025

ABSTRACT

Emotion estimation in general is a field that has been studied for a long time, and several approaches exist using machine learning. in this paper, we present an LSTM model, that processes the blend-shapes produced by the library MediaPipe, for a face detected in a live stream of a camera, to estimate the main emotion from the facial expressions, this model is trained on the FER2013 dataset and delivers a result of 71% accuracy and 62% fl-score which meets the accuracy benchmark of the FER2013 dataset, with significantly reduced computation costs. https://github.com/Samir-atra/Emotion_estimation_from_video_footage_with_LSTM_ML_algorithm

Keywords Emotion estimation · Computer vision · Social robotics

1 Introduction

In social robotics, many parts of the robot system get integrated into a single system, usually the more complex the system is, the better it will be at performing its tasks, which are more commonly focused on elderly care or patient care. building the subsystems of the robot in addition to the skeletal and motion structure, are the interaction subsystems and senses such as the speech, vision, and hearing functions, and since human emotion is a big part of any interaction between two or more humans, then it is important for the robot to understand the human emotion, and it happens through many senses such as understanding and comprehending speech, voice tone, facial expressions, body pose, and hand gesture when exists, and some times the gaze direction and face orientation could help in estimating the emotion of a person. In this work, we focus solely on emotion estimation from facial expressions.

Emotion estimation from facial expressions, more commonly referred to as Facial Emotion Recognition (FER), models could be built in three steps, that are 1) face detection: in a camera stream or a photo, and localizing the boundaries of the face, 2) landmark detection: extracting features from the face previously detected, by setting key points as a vector for the most relevant parts of the face, such as the tip of the nose and the ends of the mouth in two dimensional or three-dimensional coordinates, 3) expression classification: which comes after extracting the features by inputting them into a classification model, to predict the emotion represented by the expression shown on the face.

The system we are presenting for emotion estimation is to be implemented as a feedback system, to help drive a conversation or an interaction between a human and a robot, where it enables the robot to detect in real-time the reaction of the human through their facial expression, and based on that could change the topic or the approach of a conversation, or possibly an action it is taking.

The scope of this research was limited by the availability of data and computation resources, and to work with that the task of emotion classification was narrowed to three classes, instead of the possibility of classifying eight of them[1], and those three classes with labels happy, sad, for corresponding emotions and label unknown for every other emotion.

To start building the system will be building the full data processing pipeline, beginning from loading, cleaning, augmenting, feature extracting, and visualizing the data and then it gets inputted into the model for training. and

^{*}Work done while being a student at HAN

for the inference use of the model will integrate the trained model to the Gaze Project¹ which will do the face detection/localization and feature extraction parts, then the trained model will be classifying the face expression depending on the features.

Some considerations to be noted are, that the dataset used FER2013[2], is a challenge-dataset i.e. it is not built and optimized for research so it includes a group of non-relevant images such as animations and covered faces which does not give useful information to the model to learn from. also, work will be done on a laptop Graphics NVIDIA GeForce RTX 3050 Mobile. to train and test the model on.

This research is a proof of concept (POC), for the emotion estimation for embedded systems in social robotics full project, the main contributions this research is making are:

- 1. Building a small and cost-efficient model while taking into consideration the spatial and temporal aspects of the facial expression.
- 2. Using the MediaPipe library to localize the face and exact features for it, and using the Blendshapes[3] as features which is rarely the case when it comes to FER applications.

The rest of the paper is organized as follows: section2 Related Work, section3 Methods, section4 Ablation Studies, section5 Results, and section6 Conclusion.

2 Related Work

Facial emotion recognition is an application that has been developed and researched, for a long time[1] and many methods for face detection and tracking such as the Haar cascade algorithm, and another for feature extraction were proposed such as using the classical algorithms HOG, SIFT[4][5], or more recently using the CNN[6] and transformerencoders to detect the face and extract the features, and the classification models also using CNNs, RNNs or a combination[6][7].

2.1 Feature extraction

Feature extraction can be considered as, a special kind of data dimensionality reduction, of which the goal is to find a subset of informative variables based on image data[8]. the most common approach for extracting quality features from images of faces happens by estimating landmarks on the face, one approach to do that is finding the tip of the nose, and then segmenting it, by cropping the sphere centred at this tip, and then finding the eyebrows, mouth corners, eye corners and possibly many others on the segmented sphere, and by measuring the distance between the landmarks and changes in their positions, could understand the expression being shown on the face[9].

2.1.1 History

Looking back on the evolution of FER, one of the methods used to extract features is the Histogram of Oriented Gradients (HOG)[4], and it works by determining whether each pixel is an edge or not, the edge direction, and the magnitude of the edge. another method is the Scale-Invariant Feature Transform (SIFT)[5], which works in four parts these are: 1) construction of the scale space, 2) crucial point localization, 3) orientation assignment, and 4) assigning uniquely maintained fingerprint.

More recently, feature extraction is done using Deep neural networks[6][7], such as fully connected networks (FCN) or convolutional neural networks (CNN), and that is by building a model that consists of a number of layers and hyperparameters, determined by the complexity of the data, and the performance quality required, besides the computation resources the model meant to be implemented on.

2.1.2 MediaPipe

In this work, we used the MediaPipe[3][10] library for facial recognition and tracking, and for feature extraction, in other works such as in[11] which suggested two approaches for feature extraction, in one of them, used MediaPipe to extract features from the images, and found that using this library is better than using Deep learning networks such as CNNs, for feature extraction of all 8 emotions exists in the dataset used in that work, in the case of slight face expressions, and the more intense the face expressions get, the smaller the performance gap between the two approaches, until in the most intense case CNNs show better results in the majority of the emotions.

¹https://gitlab.com/Hoog-V/gaze/-/tree/main/gaze_estimator_python_rpi?ref_type=heads



Figure 1: Image annotated with MediaPipe

Another work is[12], where it compares MediaPipe to OpenFace, so when it comes to the landmark detection problem OpenFace[13] uses OpenCV and detects 68 facial landmarks, while MediaPipe uses TensorFlow[14], to detect 468 landmarks which are arranged in fixed quads and represented by their coordinates (x,y,z). Figure 1 shows an image annotated with the landmarks used to detect the expressions on it.

2.1.3 Blendshapes

The features to be used, are the second output from the MediaPipe Face landmark detection task model, which is the blendshapes of the frame, which are an approximate semantic parametrization and a simple linear model of facial expression[15], that was first originated in the industry before academic research and was prevalent in computer graphics. Although the blendshape technique is conceptually simple, developing a blendshape face model is a large and labor intensive effort, To express a complete range of realistic expressions, so one face might require more than 600 blendshape to be sufficiently expressed.

A single blendshape construction, was originally guided by the facial action coding system (FACS)[16], where using that system can manually code all facial displays, which are referred to as action units and more than 7000 combinations have been observed. FACS action units are the smallest visibly discriminable changes in facial display, and combinations of FACS action units can be used to describe emotion expressions[17] and global distinctions between positive and negative expressions.

MediaPipe uses the ARKit face blendshapes², which consists of 52 blendshapes that describe the face parts and expressions with probability scores in the range (0-1), showing the existence of the specific blendshape, shown in Figure2 where it shows the blendshapes histogram for the image in Figure1

2.2 Emotion classification

For emotion classification, there are many approaches that has been used in research, such as Support Vector Machines (SVM)[5] where it was implemented with linear as well as radial basis function in addition to stochastic gradient descent (SGD) classifier, which obtained results around 95% accuracy on the Radboud faces Database (RaFD)[18], and in[11] presented a comparison between two approaches, that are a MediaPipe-SVM and a CNN-LSTM where the MediaPipe and the CNN are for feature extraction and the SVM and LSTM[19] are for classification, and by experiment showed that the quality of classification based on the quality of features, and since the features extracted by MediaPipe are superior to the ones extracted by CNN, then the classification from the first combination comes with higher quality than the second.

Another image classification is by using Convolutional Neural Networks (CNN), where in[4] they used a sequential model of three convolutional layers, and a dense layer to classify the features inputted, and the method shows high accuracy on the dataset used.

Since the model being developed is meant to process videos which is a type of sequential data, then as stated in[20] much as a convolutional network is a neural network that is specialized for processing a grid of values such as an image,

²https://arkit-face-blendshapes.com/



Figure 2: Histogram for the blendshpes of the image in figure1 by MediaPipe

a recurrent neural network is a neural network that is specialized for processing a sequence of values or sequential data. so searching for the best type of RNN is a necessary step.

When working with sequential data such as videos, and needs to consider the time dependency between the images, using a recurrent neural network is the best choice for that as in[6], and using an LSTM which is a special type of RNN comes to address the vanishing gradient and the exploding gradient problems that are common in RNN training, in[21] proved by the comparison of the results, that the network including a CNN part and a RNN part delivers better accuracy, than a network with CNN layers only and that is by around 20% higher accuracy.

Another work discussing recurrent neural networks (RNN) is[7], where it suggested a CNN-LSTM architecture as a facial recognition system that could understand the spatio-temporal properties in a video.

3 Methods

In this section we will discuss our approach to building the emotion estimation system, starting with the dataset used, and the data processing techniques then building the model and optimizing it, and ending with evaluating and testing the performance of the model.

3.1 Dataset

Since the model being built is to be used in inference from video, then looking for a video dataset of different facial expressions was the first choice, and from[6] can find a group of datasets some of them are for videos such as

MMI[22][23] and AFEW7.0[24], but given the limited computation resources that are necessary for the data processing and other constraints, did not include that in the research scope, and decided to use an images dataset.

A few of the choices that were considered are MultiPIE[25], and that was because of the quality of its images due to the approach it was made by, and the suitable size of it, also considered the Radboud Faces Database RaFD[18] for the structure of the database and the quality of the distribution, another choice was AffectNet[26] once again that was because of the suitable size and the popularity of it since it is used as a benchmark to evaluate the performance of the models in many research and experiments.

For this work we used the FER2013 dataset[2] that was for its simplicity, having 48x48 greyscale images and having a useful number of images and easy to obtain being an open access dataset.

The dataset was downloaded from Kaggle³ and the data was in a .csv file representation that includes the class, image pixels array grey-scale values, and the data split, as columns and the training examples as rows.

The FER2013 dataset is structured into 7 classes that are happy, sad, angry, afraid, surprise, disgust, and neutral, and the full amount of images are split into three groups training, public test, and private test, which were used as training set, validation set and test set, respectively, and the counts of the images are as shown in Table1:

Class - Split	Full	Training	Public test	Private test
Нарру	8989	7215	895	879
Sad	6077	4830	653	594
Angry	4953	3995	467	491
Afraid	5121	4097	496	528
Surprise	4002	3171	415	416
Disgust	547	436	56	55
Neutral	6198	4965	607	626

Table 1: FER2013 dataset: classes and splits

From the table can notice that 1) The happy, sad, and neutral classes are the highest in count, while the disgust class is so rare. 2) The public test and the private test splits are fit in count to be used as the validation set and test set respectively. in the next subsections will discuss processing techniques on the dataset that will change and restructure the database in a form that best serves the model being built.

3.2 Data processing and cleaning

Using the FER2013 dataset for research while it is built for a challenge and a benchmark, comes with some requirements of data processing to make the data fit for producing an emotion estimator for a robotics application, and to do that a few steps of cleaning and processing are required.

3.2.1 Creating training data classes

The first step of data processing is to create three splits of the dataset: the training, validation, and test sets, because of the classes' count imbalance, decided not to use the full amount of available training classes' images, but include parts of them and as shown in Table2:

Class	Нарру	Sad	Angry	Afraid	Surprise	Disgust	Neutral			
Count	4000	4000	1500	1500	1500	Maximum possible	1500			
Table 2: Training set classes counts										

Choosing these numbers came because the model being built works for happy, sad, and unknown classes only, instead of the full 8 emotions' classes, so having 4000 for the main emotions will be in favour of balancing the dataset, and choosing the 1500 for most of the others because the sum of them will be 6000 in addition to the disgust class, which amounts to around 400 images, will make a reasonable count of training examples to make the unknown class to the side of the happy and sad.

³https://www.kaggle.com/c/challenges-in-representation-learning-facial-expression-recognition-challenge/data

Looking at Table1 could notice that, including bigger than these numbers of training examples from each class is possible for some of them, but including all could mean having an imbalanced data representation and a less quality distribution.

3.2.2 Readability by MediaPipe

To clean the dataset and avoid runtime errors, caused by having an image undetectable by MediaPipe, created a detection program and ran all the images of the training set through mediapipe, and excluded the undetectable ones from the dataset.

This process resulted in diagnosing the following numbers of each class as in Table3.

Class	Нарру	Sad	Angry	Afraid	Surprise	Disgust	Neutral		
Counts	377	835	597	616	239	72	261		
Table 3: MediaPipe undetectable images counts									

As described, MediaPipe detects the main parts of the face and estimates the landmark vector for each image, having a part of the face covered will make this process impossible so images such as the ones in Figure3 are examples for undetectable images.



Figure 3: Images undetectable by MediaPipe sample

3.2.3 Indexing the test set

This step came after a few experiments resulted in finding an unexplainable high error rate in the model performance, and as a tracking method added a column to the test set which is represented in a .csv file, that assigns a number for each image in the FER2013 dataset, and that number stays with the image through all the steps of data processing, and get transferred to the blendshapes of the image when the blendshapes dataset get created.

This step helps to track the image with a high error, when the model gets evaluated on the blendshapes of the test set, and be able to visualize the image or create different types of plots to check for patterns in the images causing the error.

3.2.4 Augmenting the training set

To increase the number of training images and the variety of poses and as a result, improve the generalization of the model decided to augment the data by adding new examples[27] with a random transformation applied to it, the augmentation techniques that were used are:

- 1. Random Horizontal Flip.
- 2. Random Rotation by 0.2 x 180 degrees in the clockwise direction, and counter-clockwise direction.

Choosing these techniques comes based on the use case of the model, and since it is being built for a social robot, having a flip will enable the robot to understand reflections, and provide a valid face image but with a different camera angle. while augmenting with the Random rotation would provide robustness to the model predictability, and this is because a tilted or rotated head is a common position for a human, so for the robot to be able to detect that face and classify the expressions shown on it is so helpful and important in real life.

Running the images of the dataset through the augmentation process results in a new extended dataset with more than 20000 images from the three classes described earlier.

3.2.5 Blendshapes dataset

Using MediaPipe to detect faces in a video stream, takes many preparations and processing steps for the data before it gets inputted into the classification model, and to build a model that is able to be connected to a face detection program, it needs to be trained on receiving the same output from that program, and for that processing the full dataset into a blendshapes dataset and use it to train the model.

3.3 Model architecture and training

The classification model was built fully from the LSTM layer[19] except one last Dense layer with Softmax activation to give a classification score from the last layer output.

The decision to use the LSTM layer to build the model was made for the following reasons:

- 1. The model being built is to be integrated to a video stream and classify the faces in it, which means it is a time-series data application and that means the output generated for each input to the model has a dependency on the time of its occurrence, so using one type of recurrent neural network (RNN)[20] is intuitive. Because it considers the past states besides the current input.
- 2. Another reason for this choice is that LSTM is a type of gated recurrent unit (GRU)[28], which means one of its advantages is having an inside gate that keeps out the non-relevant features from changing the loss in the model and only consider the critical ones.
- 3. LSTM stands for Long-Short Term Memory, which reflects one advantage for this type of network on the other types of recurrent neural networks, which is the ability for it to keep track for long-term dependencies well, and not only short-term dependencies, as the case with RNN[29] and GRU networks.

After choosing the type of the layer i.e. LSTM, to be mainly used, start with building the architecture of the model, and since there are many conditions to be met, that are accuracy, recall, precision, latency, and model size, and after a few experiments to find the best parameters, decided that using Keras tuner[30] as an architecture search framework to find the best values for the hyperparameters, would be more efficient. and setting the search space came based on the experiments and intuition⁴.

The final model was trained for 5000 epochs which took around two days to finish, using the batch size of 128 and the architecture was as in Figure 4, and the optimizer used was AdamW[31][32] with learning rate of 1.09e-06, global clipnorm of 1, and amsgrad was set. callbacks used were checkpoint callback and early-stopping callback.

-	r			T	ſ]		1	Г				Т	ſ				Т]		T	
lstm_input Inp	utLayer float32	-	lstm	LST	M float32		lstm_1	LS'	TM f	float32	-	lstm_2	LST sel	FM lu	float32	-	lstm_3	LSTM selu	float32		dense	Dense softma	float32
input:	output:		inpu	t:	output:		input:		outp	put:		input:		0	utput:		input:		output:		inpu	it:	output:
[(None, 52, 1)] [(None, 52, 1)]			(None, S	52, 1)	(None, 52, 52)		(None, 52	, 52)	(None,	, 52, 16)		(None, 52	16)	(Nor	ne, 52, 48)		(None, 52	, 48)	(None, 44)		(None	, 44)	(None, 3)

Figure 4: Model structure

For the loss function the categorical crossentropy was used, and for that encoded the labels of the images as one-hot vectors, and the metrics for evaluation were the loss, categorical crossentropy, categorical accuracy, and F1-score.

3.4 Model evaluation

The process of model evaluation took place on every set of weights get produced by the checkpoint callback, and shows improved training and validation performance metrics. where around every 100 epochs a model weights set gets saved.

Each model gets evaluated by predicting the labels of the test dataset images and comparing them to the ground truth labels. and a loss, categorical crossentropy, categorical accuracy, and F1-score get calculated for the model.

⁴https://medium.com/p/1833f774051f

4 Ablation Studies

To have a better understanding of the model and the possibilities for the application being developed, we conducted a few ablation studies, to test for the best number of blendshapes to be used, the type of the model, and the loss function.

4.1 Choose relevant blendshapes

To improve the efficiency of the model, lower the computation requirement for it, and be able to deliver quality classification results with less processing time and memory usage, decided to disregard a set of the blendshapes returned from MediaPipe for the face detected in each frame of the video that will be processed by the model, which is 52 according to the ARKit face blendshapes.

This idea came after noticing and as might be obvious in Figure2 that some of the blendshapes do not have a score at all, so tested for the whole dataset and found that some of the blendshapes do not have a score for all the images in the dataset.

From that notice we had two choices to use 1) count the times a certain blendshape have a zero score and disregard it, if that was a high count. 2) set a threshold number and count how many times each of the blendshapes will have a score higher than this threshold.

After testing a few values for the first choice high counts and for the second choice thresholds, decided to take the second choice and use 0.4 as a threshold and set 100 as a high count number, so for every blendshape score if it does not go higher than 0.4 for more than 100 images in the dataset, it gets disregarded.

That choice was made for the following reasons:

- Counting the zeros and making a decision based on this criteria, would have resulted a small and non-relevant number of blendshapes being disregarded, and that would not have improved the efficiency of the model after all.
- Taking the second choice and setting a threshold higher than 0.4, would have resulted in disregarding blendshapes that have a wide range of probability, which gives quality information to the model being trained to learn from, and that will get the model smaller but cut from the accuracy and performance.
- Using a higher count than 100, would have resulted in a significant cut down in the numbers of the blendshapes being used, and that would make it more difficult for the model to learn.

These reasons listed were also proven by experimenting i.e. training models on a different data collection, and testing its performance and size, then making the decision for every number mentioned.

That process resulted in using a count of 27 blendshapes out of the 52, which will decrease the size of the model while keeping the performance metrics such as the accuracy and f1-score the same.

For further clarification, Figure 5 shows the scatter plot for the 52 blendshapes and how many of them do not have high values which result in low information that could be neglected.

4.2 Dense neural network model

To get a faster model and shorter prediction time (latency) than of the LSTM network, decided to build another network that is less complex and for that chose the fully connected (Dense) layer as the only layer in the whole model. after a few optimization experiments, the metrics showed high performance from the network, but when it was integrated into the system and received the camera stream, it did not give acceptable stability in the results and kept oscillating between the classes without changing the view in the camera, besides not giving sufficiently accurate predictions.

4.3 CCE and MSE loss function

In the earlier stages of the research used the Mean Squared error[33] loss function for training the LSTM model, but because the MSE function works by calculating the distance between the prediction and the ground truth, instead of calculating the probability difference, which is fit for regression applications, had to change the loss function after a certain point.

$$MSE = \frac{1}{n} \sum_{i=1}^{n} (y_i - \hat{y}_i)^2$$
(1)



Figure 5: Blendshapes scatter plot for a sample image

In comparison to the cross-entropy, which finds the difference between the probability of having the correct category assigned to the face, and the ground truth label. which fits for classification applications. and since the application includes three classes then using the Categorical Cross-Entropy was the choice.

$$CCE = -\sum_{i=1}^{N} y_i \log(\hat{y}_i) \tag{2}$$

for possible further improvement, suggest using the Categorical Focal Cross-Entropy[34], which will take into consideration the class imbalance in the dataset. although it was originally built for object detection, experiments showed that it also improves the classification models' results.

$$FL = -\alpha (1 - \hat{y})^{\gamma} CCE(y, \hat{y}) \tag{3}$$

5 Results

The results that were obtained from the model were close to the classification benchmark of the dataset⁵ with no extra training data and using a non-transformer neural network. The best results obtained from the model on the test set were as follows:

- Loss = 0.6238
- Categorical crossentropy = 0.6235
- Categorical accuracy = 0.7199

⁵https://paperswithcode.com/sota/facial-expression-recognition-on-fer2013

• F1-score = 0.6298

and the confusion matrix for the model as in Table4 shows the classes as 0:happy, 1:unknown, and 2:sad.

	Predicted								
Ground truth	Class	0	1	2					
	0	251	35	5					
	1	110	850	150					
	2	21	140	84					
Table 4: Confusion matrix									

Experimenting with different architectures, hyperparameters and training for a longer time did not give improved metrics, but the confusion matrix keeps oscillating between an improvement in the happiness class classification one time, and for the sadness class classification another and the one above is biased to the happiness class. and for that considered assigning class weights for each one, to direct the model focus to the happy and sad classes to pay them more attention and give less weight to the unknown class which is easily classified in all experiments. but due to some limitations did not experiment with class weights.

When the model gets integrated to the system camera stream like in the demo video⁶ it indeed gives a stable output for happiness and sadness passing by the unknown class when the face look changes, and for all other facial poses the unknown class is the output.

6 Conclusion

We worked on building a model to detect faces in a camera video stream, then extract blendshapes as features from the face, and classify the expressions on the face to the emotion that it represents. and proved that a 4-layer LSTM model is a suitable architecture to classify the blendshapes of the faces in the frames of the camera stream, showing good stability with respect to time. and note that it shows good results in the video stream even after the model gets trained on an image dataset. besides the model delivers results equal to the benchmark of the dataset with no loss in accuracy through the pipeline of feature extraction and classification, which saves a lot of memory and computation in comparison to other available methods.

For future work, will be using an improved dataset and computation and work on delivering a smaller model in size and higher in performance.

Acknowledgments

Thanks to professor Marijn Jongerden, Jeroen Veen and Dixon Devasia for their help and guidance along the way, and thanks to Victor Hogeweij for his contribution to the Gaze project. Also, thanks Bhupinder Kaur, An Le, and Muhammad Reza for their insight and support at the early stages of the work.

References

- [1] Paul Ekman and Wallace V. Friesen. Unmasking the Face: A Guide to Recognizing Emotions from Facial Clues, volume 10. Ishk, 2003.
- [2] Ian J Goodfellow, Dumitru Erhan, Pierre Luc Carrier, Aaron Courville, Mehdi Mirza, Ben Hamner, Will Cukierski, Yichuan Tang, David Thaler, Dong-Hyun Lee, et al. Challenges in representation learning: A report on three machine learning contests. In *Neural information processing: 20th international conference, ICONIP 2013, daegu, korea, november 3-7, 2013. Proceedings, Part III 20*, pages 117–124. Springer, 2013.
- [3] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Yong, Juhyun Lee, et al. Mediapipe: A framework for perceiving and processing reality. In *Third workshop on computer vision for AR/VR at IEEE computer vision and pattern recognition (CVPR)*, volume 2019, 2019.
- [4] Chahak Gautam and KR Seeja. Facial emotion recognition using handcrafted features and cnn. *Procedia Computer Science*, 218:1295–1303, 2023.

⁶https://youtu.be/RdcjA9ScBmI

- [5] Pranav Kumar, SL Happy, and Aurobinda Routray. A real-time robust facial expression recognition system using hog features. In 2016 International Conference on Computing, Analytics and Security Trends (CAST), pages 289–293. IEEE, 2016.
- [6] Shan Li and Weihong Deng. Deep facial expression recognition: A survey. *IEEE transactions on affective computing*, 13(3):1195–1215, 2020.
- [7] Sze Chit Leong, Yuk Ming Tang, Chung Hin Lai, and CKM Lee. Facial expression and body gesture emotion recognition: A systematic review on the use of visual data in affective computing. *Computer Science Review*, 48:100545, 2023.
- [8] Michael Egmont-Petersen, Dick de Ridder, and Heinz Handels. Image processing with neural networks—a review. *Pattern recognition*, 35(10):2279–2301, 2002.
- [9] Akhmedov Farkhod, Akmalbek Bobomirzaevich Abdusalomov, Mukhriddin Mukhiddinov, and Young-Im Cho. Development of real-time landmark-based emotion recognition cnn for masked faces. *Sensors*, 22(22):8704, 2022.
- [10] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, et al. Mediapipe: A framework for building perception pipelines. arXiv preprint arXiv:1906.08172, 2019.
- [11] Carmen Bisogni, Lucia Cimmino, Maria De Marsico, Fei Hao, and Fabio Narducci. Emotion recognition at a distance: The robustness of machine learning based on hand-crafted facial features vs deep learning models. *Image and Vision Computing*, 136:104724, 2023.
- [12] Alexander V Savin, Victoria A Sablina, and Michael B Nikiforov. Comparison of facial landmark detection methods for micro-expressions analysis. In 2021 10th Mediterranean Conference on Embedded Computing (MECO), pages 1–4. IEEE, 2021.
- [13] Tadas Baltrušaitis, Peter Robinson, and Louis-Philippe Morency. Openface: An open source facial behavior analysis toolkit. In 2016 IEEE Winter Conference on Applications of Computer Vision (WACV), pages 1–10, 2016.
- [14] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. Tensorflow: A system for large-scale machine learning, 2016.
- [15] John P Lewis, Ken Anjyo, Taehyun Rhee, Mengjie Zhang, Frederic H Pighin, and Zhigang Deng. Practice and theory of blendshape facial models. *Eurographics (State of the Art Reports)*, 1(8):2, 2014.
- [16] Paul Ekman and Erika L Rosenberg. What the face reveals: Basic and applied studies of spontaneous expression using the Facial Action Coding System (FACS). Oxford University Press, USA, 1997.
- [17] Paul Ekman. Facial expression and emotion. American psychologist, 48(4):384, 1993.
- [18] Oliver Langner, Ron Dotsch, Gijsbert Bijlstra, Daniel HJ Wigboldus, Skyler T Hawk, and AD Van Knippenberg. Presentation and validation of the radboud faces database. *Cognition and emotion*, 24(8):1377–1388, 2010.
- [19] Neural Computation. Long short-term memory. Neural Comput, 9:1735–1780, 2016.
- [20] Ian Goodfellow. Deep learning, 2016.
- [21] Neha Jain, Shishir Kumar, Amit Kumar, Pourya Shamsolmoali, and Masoumeh Zareapoor. Hybrid deep neural networks for face emotion recognition. *Pattern Recognition Letters*, 115:101–106, 2018.
- [22] Maja Pantic, Michel Valstar, Ron Rademaker, and Ludo Maat. Web-based database for facial expression analysis. In 2005 IEEE international conference on multimedia and Expo, pages 5–pp. IEEE, 2005.
- [23] Michel Valstar, Maja Pantic, et al. Induced disgust, happiness and surprise: an addition to the mmi facial expression database. In Proc. 3rd Intern. Workshop on EMOTION (satellite of LREC): Corpora for Research on Emotion and Affect, volume 10, page 65. Paris, France., 2010.
- [24] Abhinav Dhall, Roland Goecke, Shreya Ghosh, Jyoti Joshi, Jesse Hoey, and Tom Gedeon. From individual to group-level emotion recognition: Emotiw 5.0. In *Proceedings of the 19th ACM international conference on multimodal interaction*, pages 524–528, 2017.
- [25] Ralph Gross, Iain Matthews, Jeffrey Cohn, Takeo Kanade, and Simon Baker. Multi-pie. Image and vision computing, 28(5):807–813, 2010.
- [26] Ali Mollahosseini, Behzad Hasani, and Mohammad H Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing*, 10(1):18–31, 2017.
- [27] Ajinkya More. Survey of resampling techniques for improving classification performance in unbalanced datasets. *arXiv preprint arXiv:1608.06048*, 2016.

- [28] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [29] Larry R Medsker, Lakhmi Jain, et al. Recurrent neural networks. Design and Applications, 5(64-67):2, 2001.
- [30] Tom O'Malley, Elie Bursztein, James Long, François Chollet, Haifeng Jin, Luca Invernizzi, et al. Kerastuner. https://github.com/keras-team/keras-tuner, 2019.
- [31] Diederik P Kingma. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- [32] I Loshchilov. Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101, 2017.
- [33] Mean Squared Error, pages 337–339. Springer New York, New York, NY, 2008.
- [34] T Lin. Focal loss for dense object detection. arXiv preprint arXiv:1708.02002, 2017.