
Citation:

Ruiz-Garcia, A and Elshaw, M and Altahhan, A and Palade, V (2018) A hybrid deep learning neural approach for emotion recognition from facial expressions for socially assistive robots. *Neural Computing and Applications*, 29 (7). pp. 359-373. ISSN 0941-0643 DOI: <https://doi.org/10.1007/s00521-018-3358-8>

Link to Leeds Beckett Repository record:

<https://eprints.leedsbeckett.ac.uk/id/eprint/4939/>

Document Version:

Article (Accepted Version)

The aim of the Leeds Beckett Repository is to provide open access to our research, as required by funder policies and permitted by publishers and copyright law.

The Leeds Beckett repository holds a wide range of publications, each of which has been checked for copyright and the relevant embargo period has been applied by the Research Services team.

We operate on a standard take-down policy. If you are the author or publisher of an output and you would like it removed from the repository, please [contact us](#) and we will investigate on a case-by-case basis.

Each thesis in the repository has been cleared where necessary by the author for third party copyright. If you would like a thesis to be removed from the repository or believe there is an issue with copyright, please contact us on openaccess@leedsbeckett.ac.uk and we will investigate on a case-by-case basis.

A Hybrid Deep Learning Neural Approach for Emotion Recognition from Facial Expressions for Socially Assistive Robots

Ariel Ruiz-Garcia · Mark Elshaw ·
Abdulrahman Altahhan · Vasile Palade

Received: date / Accepted: date

Abstract We have recently seen significant advancements in the development of robotic machines that are designed to assist people with their daily lives. Socially assistive robots are now able to perform a number of tasks autonomously and without human supervision. However, if these robots are to be accepted by human users, there is a need to focus on the form of human-robot interaction that is seen as acceptable by such users. In this paper, we extend our previous work, originally presented in Ruiz-Garcia et al. [1], to provide emotion recognition from human facial expressions for application on a real-time robot. We expand on previous work by presenting a new hybrid deep learning emotion recognition model and preliminary results using this model on real-time emotion recognition performed by our humanoid robot. The hybrid emotion recognition model combines a Deep Convolutional Neural Network (CNN) for self-learned feature extraction and a Support Vector Machine (SVM) for emotion classification. Compared to more complex approaches that use more layers in the convolutional model, this hybrid deep learning model produces state-of-the-art classification rate of 96.26%, when tested on the Karolinska Directed Emotional Faces dataset [2], and offers similar performance on unseen data when tested on the Extended Cohn-Kanade dataset [3]. This architecture also takes advantage of Batch Normalization [4] for fast learning from a smaller number of training samples. A

A. Ruiz-Garcia (✉) · M. Elshaw · A. Altahhan · V. Palade
School of Computing, Electronics and Mathematics
Faculty of Engineering, Environment and Computing
Coventry University, Coventry, UK
E-mail: ariel.ruizgarcia@coventry.ac.uk

M. Elshaw
E-mail: mark.elshaw@coventry.ac.uk

A. Altahhan
E-mail: abdulrahman.altahhan@coventry.ac.uk

V. Palade
E-mail: vasile.palade@coventry.ac.uk

comparison between Gabor filters and CNN for feature extraction, and between SVM and Multilayer Perceptron for classification is also provided.

Keywords Deep Convolutional Neural Networks · Emotion Recognition · Gabor filter · Socially Assistive Robots · Support Vector Machine

1 Introduction

Roboticians have long anticipated and foreseen the arrival of intelligent machines capable of illustrating human behaviours without human assistance. Social robots are now able to perform a number of tasks autonomously, which led to a steady introduction of social robots into society. The kinds of social robot applications that are becoming available include those that offer therapy to children with autism [6], exercise coaches [7] and social robots that aid elder people with dementia [8], [9]. However, these machines are still unable to effectively interact with users in a human-like manner, which is an issue that has proved to be difficult to overcome. Similar to the way human-human interaction (HHI) can be inhibited by the lack of initiative from one of the parties, human-robot interaction (HRI) can fail if there is limited or no engagement by the robot or the human user. Though, in order for the robot to take initiative, it has to be able to determine what actions to execute and to what degree. Moreover, it has to adapt its behaviours according to the user's responses. In this work, we address the first stage of an automated empathic behaviour system, which is the recognition of human emotions. We extend the work published in [1] by introducing a new hybrid architecture that combines a Deep Convolutional Neural Network, for self-learned feature extraction and representation, with a Support Vector Machine for emotion recognition. This hybrid architecture offers novelty over similar approaches in that it has a simplified configuration, less hyperparameters, and is relatively faster to train, due to the use of Batch Normalization (BN) [4] and its reduced number of layers, compared to the current state-of-the-art architectures employing CNNs. This architecture is also able to learn from smaller amounts of data than many of the similar state-of-the-art approaches.

As presented in [1], Gabor filters are common image pre-processing methods in the field of emotion recognition due to their ability to extract salient features. In the case of facial expression images these are able to highlight areas around the mouth, eyes, and eyebrows, all of which play an important role in the recognition of emotions. Considering the popularity and success of CNNs in image classification tasks, this work compares CNN to Gabor filters in terms of feature extraction for emotion recognition. However, since traditional CNN models are constrained by the efficiency of their MLP component to classify the features extracted by the convolutional layers, in this work we look at an alternative to the classifier component in an attempt to increase classification performance. In the case of Gabor filters it was shown in [1] that SVM outperform MLPs when classifying the feature vector produced by a bank of Gabor filters. Therefore, this work explores SVM as an alternative to the MLP component of traditional CNNs.

In addition, we explore the impact of moving the testing of the novel hybrid architecture from a controlled manner using images from available corpora, to a

less controlled context by applying the model on a robot in a real world environment. Hence, we present initial results on real-time emotion recognition performed by our NAO robot <<https://www.aldebaran.com/en/cool-robots/nao>> using this new hybrid emotion recognition model. Such a real robot test will identify the performance of the emotion recognition model in a real world scenario and will identify some of the factors that will need to be addressed to achieve levels of emotion recognition performance that will be required by human users.

The following section of this paper presents some background and literature review on social robots, and the current state-of-the-art of computational emotion recognition from facial expressions. Section 3 presents our improvements to the emotion recognition models originally presented in [1] and the new hybrid model. Section 4 presents the results obtained with the emotion recognition models on the Karolinska Directed Emotional Faces (KDEF) dataset [2] and the results of the best performing model from Section 3 when tested on the Extended Cohn-Kanade (CK+) dataset [3]. Section 5 describes the preliminary emotion recognition experiments carried out with a NAO robot. Section 6 presents the conclusions of the work presented in this paper along with the future direction of this research. Section 6 also highlights the importance of training robots with realistic data obtained in out-of-the-lab environments. This is followed by a list of references.

2 Background and Literature Review

The introduction of social robots into society is inevitable and the transition into a lifestyle that constitutes interacting with robots on a daily basis can be difficult if these machines do not fulfil user's interactive expectations. Therefore, social robots need to be equipped with the intelligence and skills necessary to build social, and to some extent intimate, relationships with end users. Consequently, our work encompasses the development of artificial mechanisms intended to endow social robots with the skills required to succeed in building social relations with human users. The capacity to build this social engagement between the robot and the human user will be particularly significant when considering scenarios where the interactions will be long-term. For instance, when the robot is a companion and a helper with the goal of ensuring the senior citizen can remain in their own home as long as possible.

In this work we make progress towards the development of a robot that is able to recognise emotions through the users' facial expressions. Although in this paper we focus on the goal of developing a model that is able to perform emotion recognition from facial expression, this is only the first step of our overall goal to create an empathic robot. Such a robot will not only recognise emotions but also automatically and autonomously produce and associate responses to specific emotional states. We begin by targeting the most essential characteristic of an empathic robot: the ability to recognise human emotions. This section studies existing advancements in the field of social robotics and existing machine learning approaches to perform emotion recognition.

2.1 Socially Assistive Robots

Although the concept of autonomous robots that can learn was introduced decades ago, we are just now starting to witness their introduction into society. The technological limitations that were once the main obstacle of intelligent machines are no longer an impediment. The growing trend for the use of social assistive robots can be attributed in part to the numerous benefits that these machines can offer to individuals and society. For example, scientists in Japan are focusing their efforts in developing socially assistive robots that can serve as companions for the elderly. Moreover, studies show that elderly citizens in Japan prefer in certain cases to obtain assistance from robotic caretakers rather than their human counter-parts [10]. Although this can be interpreted as a worry for some, in a place where the population is in decline and ageing rapidly, citizens can benefit from social robots by employing them as companions and caretakers for the elderly.

Other uses in the health and social care sector include the use of social robots to assist care receivers and caregivers. For example, social robots can be used to help with the coaching of care receivers to eat healthy and take exercise [11]. [7] have created a robot that encourages senior citizens to take exercise by recognising their arm movements, coaching behaviours that plan the exercise to be performances and spoken interactions with the participants. However, this robot fails to consider the mood of the user that could offer a great deal of information on why, for example, the user is not exercising to their full potential.

Social robots can also be employed by the wider spectrum domains in society. In the educational sector, social robots can be used as personal tutors and teachers and can assist in getting students to engage and learn in a proactive way. Castellano et al. [12] have identified a number of crucial points for the success of empathic robotic tutors, including the need for mechanisms for creating social bonds even if not all the features of the robot are anthropomorphic. In fact, existing studies show that social robots can increase children's interest in engineering, increase engagement in learning experiences, and improve language skills development [13]. [14] have used a socially assistive robot called Tega to help children learn a second language. In the learning session, Tega and the child learn Spanish from a virtual teacher agent in the form of a Toucan. This approach does offer emotion recognition from faces by using the Affectiva [15] system and emotion production using the behaviour and speech of the robot. Nevertheless, these emotions rely on valance to improve the engagement of the child and so the emotions considered are limited to positive, neutral and negative, which could mean that certain subtle but significant emotion shifts might be lost.

Studies have shown that social robots can be a useful resource for social skills and communication therapies for children with autism [16]. Rabbitt et al. [17] argue that parents view social robots as a very acceptable form of treatment for children with disruptive behaviour problems. As a result of the positive feedback obtained from existing case studies, researchers continue to develop frameworks designed for robots to engage in social interactions with potential users. The KSERA project is an example of this and it is specially designed to assist the elderly with conditions such as Chronic Obstructive Pulmonary Disease [18]. The KSERA robot, based on a combination of neural network learning approaches, offers activities to aid the elderly such as person

recognition, speech recognition and production and navigation. Similarly, GeriJoy, a virtual care companion, offers wellness coaching, therapeutic programs, reminders, safety supervision, companionship and care for the elderly [19]. However, unlike the other approaches considered in this subsection, this system does not rely on machine learning but is controlled by a human operator.

As social robots start to become more commonplace, the need to provide these machines with the necessary skills to build long-term social relations becomes more evident. Leite et al. [20] suggest that empathy supports the creation and development of social bonds between people and that people respond better to robots whom behave emphatically towards them. According to de Graaf et al. [21], people interact with social robots in the same way that humans interact with each other. This implies that robots need to adequately illustrate signs of intelligence and self-awareness to some extent in order to develop long-term relationships with users. Nonetheless in order to develop an empathic robot we need to first address the focus of the research in this paper, the development of a model for emotion recognition that can perform in a real-world environment. The next section reviews current advancements in the field of computational intelligence, existing state-of-the-art emotion recognition models.

2.2 Machine Learning Approaches for Emotional Face Recognition

As stated above, to make progress towards an empathic robot we need to address the first step which is emotion recognition. Hence, we will here explore the current state-of-the-art in emotion recognition.

According to Duffy et al. [22] building machines that can be as intelligent and versatile as humans, and with the ability to socialise and interact as if they were humans themselves, requires employing the human frame of reference to a certain extent [22]. Artificial neural networks (ANN) are computational methods intended to model, to some degree, the way the human brain works and can be used to classify facial expression images as a given emotion [23], [24]. Variants of this method have also succeeded in classifying facial expression images e.g. Self-organising Maps [25], [26], Support Vector Machines (SVM) [27], [28], Convolutional Neural Networks (CNN) [29], amongst others. In this work we first discuss and compare the performance of SVM, and Multilayer Perceptron (MLP) Neural Networks to classify facial expression images obtained from the KDEF dataset [2] as the following emotions: sad, surprised, neutral, happy, fear, disgust, and angry. These models rely on hard-coded feature selection using Gabor filters. In what is the main claim of this paper, we then explored an alternative approach that considered the above by replacing the Gabor filters with an approach that offers more self-learned feature selection by using a Deep CNN.

Considering that the performance of classifier algorithms heavily relies on the quality of the feature vector representing the image, and thus the emotional state, it is essential that the optimum image pre-processing method is applied to the images used for training. Gabor filter is one of the most popular methods in image processing due to its ability to detect edges. This process resembles the perception in the human visual system [30] and is characterised by multi-resolution and multi-orientation

properties. Chelali and Djeradi [33] have proposed an approach which relies on the magnitude vector produced by Gabor filter. The authors propose applying discrete wavelet transform (DWT) to images as a pre-processing step before being classified by MLP and Radial Basis Function networks. Chelali and Djeradi [33] obtain a peak performance of 95% accuracy on the Computer Vision database and 85% on the ORL database. Mehta and Jadhav [31] use a combination of Log-Gabor filters, PCA, and Euclidean distances on the JAFFE <<http://www.kasrl.org/jaffe.html>> dataset and obtain a performance of 93.57%. Ahsan et al. [30] also use Gabor filters and combine them with Local Transitional Patterns (LTP): the authors apply Gabor wavelet filter on images and then obtain LTP codes by comparing transition of intensity change at different levels of neighbouring pixels in different directions. The resulting feature vector is then classified with a SVM, which produces an average accuracy rate of 95% on the Cohn-Kanade (CK) [32] database.

Facial expression classification using SVM has also been done by Sohail and Bhattacharya [28]. The authors proposed a method which includes identifying 15 different feature points and measuring the Euclidean distances between these and the feature points representation of a neutral face. The authors employ a SVM for classification and obtain an average recognition rate of 92% on the JAFFE dataset and 86.33% on the CK [32] dataset. These facial expression classification results using SVMs as a classifier positively illustrate the strengths of SVMs for emotion recognition. SVMs have also been employed for other classification tasks such as face recognition problems [34][35], and face formalisation [36].

MLPs have also proven to be efficient for facial expression classification. Hewahi and Baraka [37] designed a MLP which makes use of ethnic background information to produce an accuracy rate of 83.3%. When no ethnic background information is used, the authors obtain a performance rate of 75% with the same model on the MSDEF dataset. Khashman [38] obtain 87.78% accuracy rate when using Global Pattern Averaging as an image pre-processing step and a MLP for facial expression classification.

Due to the success of CNNs in image classification problems, CNNs are becoming of interest in facial expression classification problems. [29] uses a CNN with five convolutional layers to obtain a feature vector and then feed it to a SVM for classification. The author obtains 94.4% accuracy rate on the Extended Cohn-Kanade Dataset. Similarly, Burkert et al. [39] have set a benchmark with 99.6% accuracy using a Deep CNN. The authors use a CNN with seven convolutional layers and test it on the CK dataset using a 10-fold cross-validation method.

Although the models discussed in this section offer a good degree of accuracy, it is unclear what features determine classification performance. According to Beaudry et al. [40], the eyes and eyebrows play a bigger role for the recognition of sadness and the mouth is more influential to recognise happiness. In contrast, the authors determined that a holistic processing could be called upon fear but could not determine the best approach to recognise other emotions. Identifying what facial features determine specific emotions could be of great value if applied to emotion recognition models. In this work, we try to emphasise on areas of interest such as the mouth, eyes and eyebrows to classify facial expression images. As a result, in this paper, we explore Gabor filters as an image pre-processing technique to highlight these areas within an

image. Moreover, we use CNN in our new architecture to extract features and compare them against those obtained with Gabor filters. The following section describes our feature extraction methods and classification models.

3 Methodology

This section presents the emotion recognition approach originally published in [1] and introduces an architecture that employs a CNN for automatic feature extraction and representation. The latter architecture is less complex to train, needs less data to learn, and achieves results at a similar level of performance compared to more complex and larger architectures. A visualisation of the filters learnt by the CNN and those obtained with the Gabor filters is provided. Both models use SVMs and MLPs as classifiers.

3.1 Emotional Facial Expression Corpus

In this work, we train and test our emotion recognition models on the Karolinska directed Emotional Faces [2] database. The corpus contains a set with 70 individuals: 35 males and 35 females aged between 20 and 30 years, each displaying seven different emotional expressions in five different angles. All images were taken under a controlled environment and faces were centered with a grid by positioning eyes and mouths in fixed image coordinates [2].



Fig. 1: Subject F07 from the KDEF [2] dataset, displaying seven emotions: sad, surprised, neutral, happy, fear, disgust, and angry.

During our experiments we only use front angle images; a subset containing 140 front angle images for each one of the seven emotions. In order to obtain a feature vector, we located the face and cropped irrelevant spatial features such as background. Face images were grayscaled and resized to a standard 120×110 in the case of Gabor filter and 100×100 in the case of CNN to speed up the training. Figure 1 illustrates sample face images obtained from the KDEF database.

We also test the CNN+SVM model on unseen data using the Extended Cohn-Kanade dataset [3], a popular dataset in the field of emotion recognition and face detection. This corpus includes a set of 327 sequences from 118 participants. In this work, we only use the peak frame of each sequence since it contains the most emotion-related information. The same pre-processing applied on the KDEF dataset is applied to this dataset for consistency, as shown in Figure 2. Testing on the CK+



Fig. 2: Left to right, subjects S52, S55, and S116 from the CK+ [3] dataset, displaying six emotions: surprise, happy, disgust, angry, fear and sad. Note that no participant in the dataset provided consent for *contempt* images to be distributed.

dataset allows us to compare our hybrid model to the work of [29] who also uses a combination of CNN+SVM.

3.2 Gabor Filter

Gabor filters are powerful image processing algorithms that resemble the perception in the human visual system [30] and facilitate edge detection on images. Facial expressions classification heavily relies on the shape of facial features such as the mouth, eyes, and eyebrows [40], and Gabor filters can be used to emphasise these areas. Therefore, for our first experiment we convolved our dataset with a bank of Gabor filters to obtain image representations that highlight these areas of interest. Each Gabor filter used is essentially a sinusoidal modulated by a Gaussian kernel function [33] in which orthogonal directions are represented by real components.

In [1] we used a bank of 40 filters with five scales and eight orientations, split the resulting feature vector into four and treated each sub-vector as an input sample in order to reduce overfitting. However, because the previous approach quadrupled the amount of data, in this work we only use a bank of 10 filters expanding over eight orientations and two dimensions. This is done to provide a fair comparison against the model using a CNN as a replacement for the bank of Gabor filters and using the same number of training and testing samples. Moreover, we did not observe any improvements in classification performance when using more filters, and using a smaller number of filters allowed us to craft each one to highlight areas of interest such as the mouth, eyes and eyebrows.

The Gabor filters applied also down sample the original image over a scale of four, producing a reduced feature vector. Using a feature vector with reduced dimensionality allows for faster training of the SVM model, and also produces similar results to full scale vectors. Once the Gabor filters were applied, the feature vector values were normalised in the range zero to one. Moreover, as done by Chelali and Djeradi [33], we only used the magnitude information given that it contains the most relevant information and discards the effect of noise.

We tried using a combination of real and imaginary components as a complex component, however we obtained lower classification results compared to using only the real component. Let λ represent the frequency of the sinusoidal, θn represent the orientation, and σ represent the standard deviation of the Gaussian over x and y dimensions of the sinusoidal plane; our Gabor filter applied to an image with dimensions x and y is defined as follows:

$$g_{\lambda,\theta}(x,y) = \exp \left[-\frac{1}{2} \left\{ \frac{x_{\theta n}^2}{\sigma_x^2} + \frac{y_{\theta n}^2}{\sigma_y^2} \right\} \right] \cos(2\pi * \theta n * \lambda) \quad (1)$$

where

$$x_{\theta n} = x(\sin \theta n) + y(\cos \theta n)$$

$$y_{\theta n} = x(\cos \theta n) + y(\sin \theta n)$$

After trying a number of parameters, we concluded that initialising the Gabor filter with the following parameters produces the best magnitude response vector for emotion classification: $\theta = 2\pi/3$, $\lambda = 6$, $\gamma = 0.5$, and $\sigma = 4$. This response vector is given by:

$$\|g_{\lambda,\theta}(x,y)\| = \sqrt{\Re^2\{g_{\lambda,\theta}(x,y)\} + \Im^2\{g_{\lambda,\theta}(x,y)\}} \quad (2)$$

where \Re denotes the real component of the filter and \Im the imaginary one.

3.3 Facial Expression Classification using SVM and MLP

Once the feature vector was obtained, we feed it to a SVM and a MLP to be classified as one of seven emotions: angry, disgust, fear, happy, neutral, sad, and surprised. SVMs are non-probabilistic binary classifiers known for performing notably well in image classification problems. We tested the one-vs-one and one-vs-all approaches for multi-class classification and obtained better performance with the one-vs-one approach using a linear kernels. Let b represent the bias, K be a linear kernel function, our facial expression classification is determined by:

$$f(x) = \text{sgn} \left(\sum_{i=1}^l y_i a_i [K(x_i, x) + b] \right) \quad (3)$$

where x_i is the training vector, x is the testing vector with $a_i > 0$, y_i represents Lagrange multipliers of dual optimization problem [42]. This model produced an accuracy rate of 95.58% on the testing set after training using a c value of 1000 for the SVM. We also tested on larger and smaller c values but obtained better results with 1000.

Taking into account the popularity of traditional MLPs for classification problems, we decided to compare the performance of the SVM against that of a MLP network. After considering a number of different network topologies we obtained best results with the following MLP network configuration: one input layer with 8400 neurons taking ten 28×30 filtered images as input, one hidden layer with 93 neurons, and one output layer with 7 neurons. The target values contained a one in the place of the target class and zero for the rest. The MLP uses a sigmoid activation function and was trained using Resilient Backpropagation (Rprop) to reduce the issues of falling into a local minimum caused by sigmoid activation functions. Figure 3 provides an illustration of this model. The Gabor+SVM approach is the same as in figure 3, but we replace the MLP in the figure with a SVM.




Fig. 3: Multilayer Perceptron emotion recognition model taking an image convolved with ten Gabor filters as input. Hidden layer with 93 neurons and output layer with seven. Face image extracted from the KDEF dataset [2].

In order to compare against the performance achieved with the SVM, we randomly selected 70% of the input vector as our training set and the remaining 30% for testing. The initial weights were randomly initialised and this model achieved its best performance after training for 175 epochs, with 100% percent accuracy on the training set and 93.5% on the testing set. Learning rate was set to 0.0001 and remained constant during training.

3.4 Convolutional Neural Networks




Fig. 4: Illustration of CNN+SVM and CNN+MLP emotion recognition models. CNN topology: Conv, BN, ReLU, MaxPool; Conv, BN, ReLU, MaxPool; Conv, BN, ReLU, MaxPool; Conv, BN, ReLU. MLP topology: Hidden, BN, ReLU. Face image extracted from the KDEF dataset [2]. Only one classifier, SVM or MLP, is applied at a time.

Although Gabor filters offer good image representations and allow for the detection of edges and, thus, facial features that are important for emotion recognition, Gabor filters need to be carefully designed to produce such representations. In contrast, Convolutional Neural Networks offer an alternative to prescribed methods and possess the ability to learn to extract features necessary for emotion recognition automatically. CNNs are a special type of neural networks that take advantage of spatial information in images. Similarly to Gabor filters, CNNs are biologically inspired and resemble the animal vision cortex [25]. Moreover, the convolutional layers in a CNN provide some degree of shift and deformation invariance [25]. The filters learnt by convolutional layers are similar to those produced by a bank of Gabor filters as illustrated in figure 5.

Our CNN emotion recognition model is composed of blocks of Convolutional, ReLU, MaxPooling, and BN layers [4]. We use BN since it allows for larger learning rates and faster convergence by normalizing the distribution of each input feature at every layer [4]. BN has been shown to speed up convergence by reducing internal covariate shift, that is the change in distribution of network activations caused by change in the network's parameters during training, by using the mean and variance of each minibatch to normalise activations [4]. In the case of convolutional layers this normalisation is done for each individual feature map.

The first two convolutional layers have filters of size 5×5 and the last two layers have filters of size 3×3 . All convolutional layers have a sliding window of size one and zero padding of size two. Max pooling is done with a stride of size two, kernels of size 2×2 and zero padding of size one. Using larger filters seemed unnecessary given that the size of the images is only 100×100 and bigger filters often miss relevant information. Moreover, we did not observe any improvements in classification performance using more layers, while reducing the number of layers decrease performance by three percent on average.

The output of the last block is connected to a fully connected layer, which in effect is a MLP with 100 neurons and a ReLU activation function. We also tried using a larger MLP, but saw no significant improvements. The output of this layer is also normalized with a BN layer before being classified as one of the seven classes. Figure 4 illustrates a visualization of this model when using either a MLP and or a SVM as the classifier.

The output of convolutional layers in our model can be summarized as:

$$C(x_{u,v}) = (x + a)^n = \sum_{i=-\frac{n}{2}}^{\frac{n}{2}} \sum_{j=-\frac{m}{2}}^{\frac{m}{2}} f_k(i, j) x_{u-i, v-j} \quad (4)$$

where f_k is the filter with a kernel size $n \times m$, applied to the input x . In our models n is always the same as m . The convolutional layers in the first network use 20, 40, 60, and 30 filters respectively. Every output of a convolutional layer in our models is shaped by a ReLU function. The feature vector is further reduced with pooling layers using the max operator. Furthermore, our models use a fully connected layer which in effect is a MLP. Let σ represent a ReLU activation function, then the output of the hidden layer is computed by:

$$F(x) = \sigma(W \times x) \quad (5)$$

The CNN is trained using stochastic gradient decent for 500 epochs as follows: the learning rate is set to 0.1 and dynamically adjusted down with a decay of 0.01. Let λ represent the initial learning rate, θ represent the learning rate decay, and ω the current epoch, the learning rate LR is adjusted according to:

$$LR = \frac{\lambda}{1 + (\omega \times \theta)} \quad (6)$$

During training, the output of the network is shaped by a SoftMax operator and the cross-entropy loss y is defined by:

$$y = -x_c + \log \left(\sum_j \exp(x_j) \right) \quad (7)$$

where c is the class ground-truth. The classification obtained with the network is 91.16% on the testing set. We also tested this network with a SVM as a classifier instead of a MLP. Once the CNN is trained, we remove the classification layer together with the last BN and ReLU activation layers. Classification is then performed using a SVM as follows: the entire dataset is passed through the network and the resulting downsampled feature vector is used to train the SVM. The accuracy obtained with this approach was 96.26%.

4 Emotional Recognition Models, Results and Discussion

The aim of this work is the development of a new architecture for an emotion recognition model that can allow a robotic companion to recognise human emotions in real-time and in an unconstrained environment. As an initial step, we have developed a set of emotion recognition models that employ SVMs and MLPs as classifiers using Gabor filters for feature representation.

Table 1: Overall accuracy on the KDEF dataset per class per emotion recognition models. Left to right: model using Gabor filters and MLP, model using Gabor filters and SVM, model using CNN and MLP, model using CNN and SVM.

	Gabor+MLP	Gabor+SVM	CNN+MLP	CNN+SVM
angry	96.81	90.48	88.10	95.24
disgust	91.05	97.62	88.10	95.24
fear	93.90	97.62	78.57	90.48
happy	96.34	95.24	95.24	97.62
neutral	92.41	100	97.62	100
sad	94.60	95.24	92.86	95.24
surprise	88.00	92.85	97.62	100
average	93.5%	95.58%	91.16%	96.26%

The emotion recognition model that employed Gabor filters and a MLP achieves an overall accuracy of 93.5% on the KDEF dataset. The model that applied a multi-class SVM, also using a bank of Gabor filters for feature extraction, produced a state-of-the-art accuracy rate of 95.58% also on the KDEF dataset.

Our new architecture, that only includes four convolutional layers, is trained as a single unit using mini batch stochastic gradient decent and takes advantage of BN for a faster convergence. When the classifier is a MLP, the accuracy produced by this network is 91.16%. However, when the classifier is a SVM, it produces a state-of-the-art accuracy rate of 96.26% on the KDEF dataset. Table 1 illustrates the performance of all the models on each class of the dataset and the overall accuracy on the test set. Table 2 illustrates the confusion matrix produced by both models using SVM as a classifier.

Table 2: (a): Gabor + SVM emotion recognition model confusion matrix on the test split of KDEF dataset; (b): CNN + SVM confusion matrix on test split of KDEF dataset. A: angry; D: disgust; F: fear; H: happy; N: neutral; Sa: sad; Su: surprised. Rec: recall, Prec: precision. Vertical axis: true labels. Horizontal axis: predicted labels.

	A	D	F	H	N	Sa	Su	rec		A	D	F	H	N	Sa	Su	rec
A	38	0	0	0	3	1	0	90.48	A	40	0	1	0	0	1	0	95.24
D	0	41	0	0	0	1	0	97.62	D	1	40	0	0	0	1	0	95.24
F	1	0	41	0	0	0	0	97.62	F	1	0	38	0	0	2	1	90.48
H	2	0	0	40	0	0	0	95.24	H	0	0	0	41	1	0	0	97.62
N	0	0	0	0	42	0	0	100	N	0	0	0	0	42	0	0	100
Sa	0	0	1	0	1	40	0	95.24	Sa	0	0	1	0	1	40	0	95.24
Su	2	0	1	0	0	0	39	92.85	Su	0	0	0	0	0	0	42	100
prec	88.37	100	95.35	100	91.3	95.24	100	95.58	prec	95.24	100	95.0	100	95.46	90.90	97.67	96.26

(a)

(b)

As it can be observed in Table 1 we obtained best results when using SVM as a classifier regardless of the feature extraction method. However, our CNN model slightly outperforms Gabor filters as a feature extraction method, at least for this particular dataset. One of the main differences that can be observed in Table 2 between CNN+SVM and Gabor+SVM models is that CNN+SVM classified all the *surprise* images correctly obtaining a recall of 100%, whereas the Gabor+SVM model obtained a recall of 92.85%. Another discrepancy happened with *angry* emotions, in which the CNN+SVM model classified 95.24% of them correctly compared to 90.48% by the Gabor+SVM model. The Gabor+SVM model also produced a low precision score on *angry* due to a significant number of false positives. Both models produced the same performance, including miss-classifications, on *sad* and *neutral*. The main advantage on classification performance from the CNN+SVM model over that of the Gabor+SVM model was on *angry* and *surprise*. And the main advantage of the Gabor+SVM model over the CNN+SVM model was on *Fear*, for which the CNN+SVM model obtained the lowest classification.

It is also evident that SVM produced higher classification rates compared to traditional MLP networks in the case of the dataset used for these experiments. However, we observed that the accuracy rate produced by SVM is largely dependent on the image pre-processing techniques applied to the data. We tried training the SVM with features obtained from the third convolutional layer but observed a decrease of at least ten percent in performance. Also slight changes in wavelength or orientation in the Gabor filters greatly influenced the classification performance of the SVM. Another advantage offered by CNN, compared to Gabor filters, is that they learn to extract a feature vector that represents the input image automatically, whereas Gabor filters need to be carefully crafted and optimised by means of trial and error.

It can be observed in Table 1 that although the CNN+SVM model produces a substantially higher accuracy rate than the CNN+MLP model, the ratio on misclassified labels remains almost the same. In both cases fear is the most misclassified class. There is also a correlation on similar performance on particular emotions. Both models produced equal performances on angry and disgust, and neutral and surprise.

Another observation made was that although training a CNN often requires a lengthy process it is a more automated process than hand-crafting Gabor filters. The bank of Gabor filters used in this work required a number of trial-and-error session given that slight changes in orientation or wavelength change the performance of the SVM significantly. These effects are not observable until a SVM is trained with the Gabor features. Therefore, with these observations, we conclude that Convolutional Neural Networks offer a self-learning alternative approach to Gabor filters for feature selection in the domain of emotion recognition.




Fig. 5: a - Magnitude response from Gabor filter; b - Output of the 4th convolutional layer of the CNN.

Figure 5 illustrates a side by side comparison of the image representations produced by our bank of Gabor filters and by the last layer of our convolutional network. As it can be observed the features learnt by the CNN are much simpler than those produced by the bank of Gabor filters, this in effect can be explained by the loss of information caused by pooling layers in the CNN. However, even though the features extracted from the CNN contain much less information they still retain important

information necessary for emotion recognition. Furthermore, when looking at the filters learnt by the first convolutional layer these looked relatively similar to the ones produced by the bank of Gabor filters. Essentially, the first layer of the CNN learns Gabor-like filters, which are then improved further through deeper layers. Nonetheless, the representations created by both the CNN and Gabor filters highlight facial features such as the eyes, eyebrows, and mouth, which allows us to confirm their key role in the recognition of emotions.

Table 3: CNN + SVM emotion recognition model confusion matrix on the CK+ dataset after leave-one-subject-out cross-validation. A: angry; C: contempt, D: disgust; F: fear; H: happy; Sa: sad; Su: surprised. Note that none of the images were used to train the CNN component.

	A	C	D	F	H	Sa	Su	total
A	42	1	1	0	0	1	0	93.33
C	1	16	0	1	0	0	0	88.88
D	0	0	59	0	0	0	0	100
F	0	0	0	22	2	1	0	88.0
H	0	0	0	0	69	0	0	100
Sa	3	1	0	0	0	22	2	78.57
Su	0	1	0	0	0	0	82	98.79
								95.87%

A similar hybrid model to the CNN+SVM model using five convolutional layers has been proposed by Ouellet [29]. The model proposed by the author uses a CNN with five convolutional layers and achieves a classification performance rate of 94.7% using a leave-one-subject-out cross-validation approach. In order to compare the CNN+SVM emotion recognition model proposed in this work against that of [29], we tested it on the CK+ dataset using the same leave-one-subject-out protocol. Our hybrid model achieved a classification performance of 95.87% on the CK+ dataset, as illustrated in Table 3 above. Note that the CNN element has never seen images of the CK+ dataset and was trained on a subset of the KDEF dataset as discussed above.

Our new architecture approach slightly outperforms the model proposed by [29]. With the advantages over that proposed by [29] being that the CNN component of our model was trained from scratch on a subset of 686 images extracted from the KDEF dataset, whereas the CNN component of the model proposed by the author was originally trained on 1.2 million images from ImageNet in the Large Scale Visual Recognition Challenge 2012 before being tested on the CK+ dataset. Moreover, our model could be more suitable for real-time emotion recognition given that it uses images of smaller size, 100×100 compared to 227×227 , and only has four Convolutional layers compared to five, making it faster to classify facial expression images. The fast convergence of the network, an average of 60 minutes on a quadcore processor for 500 epochs, gave us the opportunity to test different topologies and find the optimal parameters. When adding and removing layers to our model we observed

either a decrease or no significant increase in performance. We attribute the fast convergence rate to the CNN’s simplified configuration and the use of BN which allows for larger learning rates [4]. BN also eliminates the need for other heuristics, such as Dropout, that may result in a loss of information.

5 Initial Real-time Robot Emotion Recognition




Fig. 6: NAO robot used in our experiments.

Taking into account the key role played by human empathy in HHI, we believe that endowing social robots with the ability to recognise a set of emotions in an unconstrained environment will allow us to move on to the overall goal of creating an empathic robot. In this work, we have presented a set of emotion recognition models that provide state-of-the-art classification performance on the KDEF dataset. However, since the goal of our research is to move towards empathic robots, it is imperative to test the performance of these models in unconstrained environments in which an empathic robot may be used. We do not address empathy in robots in this work, but rather focus on the first stage of it, emotion recognition, and attempt to highlight the issues that the robot might find when performing real-time emotion recognition. Therefore, we have integrated our hybrid model, which produces the best accuracy rate on the KDEF dataset, within a humanoid robot. This section presents our initial experiments on real-time emotion recognition performed by a NAO robot and discusses observations made. As this is an initial experiment, the results presented are provisional and our goal is to establish how the model performs in a real world application and what issues we will need to consider when making the shift to emotion recognition using a robot.

5.1 Experimental Set-up and Methodology

This experiment was carried out using our humanoid NAO robot, a 58 cm in height robot with a number of sensors and abilities such as moving, feeling, seeing, speaking, hearing, and thinking [43]. Refer to figure 6 for an illustration. The NAO robot




Fig. 7: a: Sample image from subject F1 illustrating an angry expression, image misclassified with disgust; b: Correctly classified sample image from subject M2 illustrating a happy expression. Actual labels assigned according to participant’s feedback.

was set up in a cluttered room with average lighting conditions: lights were kept on and room has windows which allows for ever-changing natural lighting. The robot was placed on a one meter high surface in order to be able to track participant’s faces. NAO performed real-time emotion recognition on four participants: two males and two females aged between 20 and 25 years old. Participants were randomly selected and had different ethnic backgrounds, as opposed to participants in the dataset used to train the emotion recognition model who are of white ethnic background. When interacting with the robot, participants were asked to express four different emotions of their choice.

NAO was programmed to track participant’s faces and take an image of the participant once a face is detected. Participants were asked to stand within a distance of 1.5 meters away from the robot and express an emotion once NAO confirmed it detected their face: NAO fixes its head in the direction of the person standing in front and turns its eyes green to confirm it detected the user’s face. Participants were asked to perform the emotional expression in a natural way.

Since the robot does not have enough processing power, classification of the image is done off-board. The image is sent to an external laptop for processing. The participant’s face is then extracted and grayscaled. Once a face image is obtained, this is fed to our CNN+SVM hybrid emotion recognition model which returns a predicted value for the given image. When classification of the image is completed, NAO receives a command to express an appropriate response to the user’s emotional state that could in the future be replaced with learned empathic behaviour: for example, express excitement through speech and body language if the classification result was *happy* or express sympathy and support if the classification was *sad*. Given the purpose of this experiment is to highlight the issues that we will need to address in future work, we did not focus on having the robot perform the right action that would best improve the interaction process with the user. Moreover, we focused only on the recognition of emotions in the user and, thus, had the robot perform hard-coded actions. We performed four trials and asked participants to change the emotion they express each trial. Figure 7 illustrates sample images obtained by the robot and fed to our CNN+SVM emotion recognition model for classification.

Table 4: Real-time emotion recognition results: *F* and *M* denote participant’s gender; *emotion* represents emotions expressed by participant and *label* the label predicted by NAO. Last row shows the overall classification accuracy.

Participant	Trial 1		Trial 2		Trial 3		Trial 4	
	Emotion	Label	Emotion	Label	Emotion	Label	Emotion	Label
Subject F1	happy	disgust	disgust	disgust	sad	sad	neutral	neutral
Subject F2	angry	disgust	disgust	disgust	fear	fear	happy	happy
Subject M1	angry	angry	neutral	sad	surprise	surprise	sad	neutral
Subject M2	happy	happy	neutral	neutral	sad	neutral	surprised	surprised
accuracy	2/4		3/4		3/4		3/4	
	11/16							

5.2 Real-Time Emotion Recognition Results and Discussion

Performing real-time emotion recognition in unconstrained environments is a challenge difficult to overcome for social robots. One of the main obstacles faced by these machines is the ever-changing environments, which make it difficult to obtain facial expression images of similar quality to ones used to train the emotion recognition models employed by the robot. We have conducted preliminary real-time emotion recognition experiments with our humanoid robot in an attempt to obtain an understanding of the models performance on unseen, and marginally different, data. The robot used our hybrid CNN+SVM emotion recognition model which produces an accuracy rate of 96.26% on the KDEF dataset. When integrated with the robot, the CNN+SVM model produces a significantly lower average accuracy rate of 68.75%. Although testing of the model on the robot is preliminary, these findings do indicate the impact of the placing the model on a robot in an unconstrained environment and will direct our research in terms of adapting our emotion system so it achieve comparable results to testing using corpora images collected in controlled environments.

We performed four trials on four participants. Each trial consisted of each participant illustrating an emotion, not performed in previous trials, in front of the robot. Table 4 illustrates the emotions illustrated by each participant and the predicted label by the robot using our hybrid CNN+SVM emotion recognition model. The decrease in performance is attributed to the significant difference between the images obtained by the robot and those in the KDEF dataset. Moreover, the emotion recognition model has never seen images of any of the subjects before. In addition to this, in the KDEF dataset all faces are positioned in fixed coordinates. Whereas in our experiment, participants stood at varying distances from NAO, and given the difference in height the position of each participant’s face is located at different coordinates in the image space. Furthermore, the lighting conditions varied during each trial. Nonetheless, these are preliminary results intended to shift our focus onto issues that will impact the robot’s performance. As a consequence, future work will focus on performing emotion recognition with different light conditions and for faces at different angles.

One of the observations made was that NAO confused happy with disgust and sad with neutral. In the case of sad and neutral, it can be explained by the similarity between the two expressions, however in the case of happy with disgust it is difficult to come to a conclusion given that happy is not very similar with disgust. However, as it can be observed in figure 7a, the *angry* expression illustrated by participant F1 can easily be confused with *disgust*, as done by our emotion recognition model. In this particular instance, *disgust* was the class with the highest value and *angry* the one with the lowest value. Another observation made was that, although participants were asked to illustrate emotions in a natural way, they reported emphasising their facial expressions more than they normally would since they were conscious of the difficulty of recognising emotions.

Another issue that needs to be taken into account is the technological limitations of our robot, and most other robots. NAO has a pair of built-in cameras that provide a maximum resolution of 1280×960 at 30 frames per second. Better performance may be achieved with higher resolutions which may capture facial features better, such as wrinkle lines, which play a key role in the classification of emotions. Although not part of the experiments, we have also observed that people react positively to the empathic behaviours illustrated by NAO after identifying an emotion. This highlights the importance of not only being able to identify an emotion in the user, but also being able to respond with appropriate behaviour according to the identified emotional state.

The main purpose of this experiment was to highlight the issues that social robots will encounter when recognising emotions in real-time. When analysing the scenarios where the robot failed and succeeded we observed two main differences between the images correctly classified and those misclassified. First, as it can be observed in figure 7, the lighting on image 7a is much different than that of image 7b. This was also observed on most of the images where the robot failed to detect the right emotion. Second, two of the misclassified images show the participant's faces with a slight angle and tilt. Both of these observations can be justified by the fact that the model was trained on a dataset that contains images of very similar quality. The authors of the KDEF dataset also explain that they centred the faces with a grid. We hypothesize that the robot would greatly benefit from an emotion recognition model trained with much larger datasets containing a wide variety of images taken in unconstrained environments.

In this work, we have presented preliminary results on real-time emotion recognition performed by our humanoid robot in an unconstrained environment. The purpose of this experiment is to obtain an understanding of the performance of emotion recognition models in such environments and highlight the issues affecting performance to be addressed in future work. Our real-time emotion recognition results showed that NAO successfully recognised participant's emotional states 11 out of 16 times. The main advantage offered over other systems that perform real-emotion recognition is that our model recognises seven different emotions compared to the traditional positive, neutral, and negative emotional states [15]. Moreover, these results allowed us to conclude that emotion recognition in unconstrained environments is possible, though there are some issues that need to be considered in future work, such as illumination and face angle invariance. Furthermore, these results allow us to progress towards

automated emotion recognition in social robots, which will, in the near future, allow us to create a system that does not only allow our robot to identify seven different emotional states in the user, but also empathise with them.

6 Conclusions and Future Work

In this work, we have presented a set of emotion recognition models trained and tested on the KDEF dataset. Our best model was a combination of a CNN and a SVM and it produced a state-of-the-art performance rate on the KDEF dataset and comparable results to larger models on the CK+ dataset. When tested on the CK+ dataset, note that none of the CK+ images were used to train the CNN component of the model; this hybrid CNN+SVM architecture produced slightly better results than a larger model proposed by [29] and falls 3.73% short from the state-of-the-art [39]. Nonetheless, our hybrid model uses less data to train, converges relatively quickly, and has a smaller number of model parameters. This work has also showed the advantages of CNN over Gabor filters in terms of feature extraction and the advantages of SVM over MLP for feature classification, at least in terms of emotion recognition.

This hybrid architecture offers novelty over similar approaches, also employing a CNN, in its simplified configuration that requires less hyperparameters, is relatively faster to train compared to the current state-of-the-art, and is able to learn on smaller amounts of data. Moreover, this hybrid architecture offers comparable, and in some cases better, classification performance rates than larger and more complex architectures. This hybrid architecture also takes advantage of BN for a faster convergence. Note that one of the main limitations of this work is that it does not perform as well in real life scenarios, as discussed in section five, though this will be explored in future work.

In future work we will shift our focus to the Deep CNN model, given the advantage that it offers over Gabor filters for automatic feature extraction and representation. We will explore the performance of the same model when the convolution layers are pre-trained as stacked Auto-Encoders; each layer is trained to encode the input and a layer is added to decode the downsampled representation obtained by the first layer. Once all layers are trained individually, they are combined into a single model and the classification layer is added. The entire model is then trained, for the classification layer, and fine-tuned, for the convolutional layers. This method has proven to be successful in the past [41]. In addition to this, we will explore the possibility of reducing the number of layers or using a random set of patches as a representation of the image, in order to further improve training time by taking advantage of the hybrid model's simplified architecture and its ability to learn on a smaller number of training samples.

We have also provided preliminary findings on real-time emotion recognition performed by our NAO robot using the new hybrid CNN+SVM model. This experiment was carried out in an unconstrained environment and was intended to highlight the issues that a social robot may face when performing real-time emotion recognition. Letting our robot to perform real-time emotion recognition allowed us to realise the importance of training the emotion recognition models with realistic data, taken in

unconstrained environments where the robot will be used. This has been pointed out by Castellano et al. [12] in the past. As a result, we will look into the development of a dataset with images taken in real life scenarios in which a social robot could potentially be used. This will also allow us to train the robot with images that are not always perfect, do not have the same illumination conditions or have the perfect angle. This is because the robot will not always be at the same height of the user or be able to obtain an image with the face centred. Although there exist methods to correct this, they would only add complexity and delay the response of the robot. In an ideal scenario, an emotion recognition model should produce similar performance on the dataset used for training and in the environment in which it will be used, regardless of the environment conditions.

Once an emotion recognition model produces promising performance for real-time emotion recognition on the NAO robot, the next step will be to develop a model that allows a robot to learn the same emotional state within its very own system: composed of sensor and motor values. We hypothesize that applying the properties of mirror neurons, i.e. simultaneous action perception and execution, to a neural architecture would allow us to create a system that allows the robot empathise with users in real-time and with self-learned actions rather than hard-coded ones as done in this work. This model will have to take into account that the robot is constrained by its ability to demonstrate emotional states using body language and facial expressions. This will naturally be subject to the robot being used and its anthropomorphic characteristics. Finally, we will explore the inclusion of Reinforcement Learning for continuous learning within the robot and allow it to adjust its behaviours according to the user responses. We believe that this is a viable path to a more personalised and possibly intimate interaction between robots and humans.

In this work we made progress towards the development of an empathic robot, a robot with the ability to (i) recognise human emotions through facial expressions, (ii) illustrate emotional states itself, and (iii) automatically and autonomously produce and associate responses to specific emotional states. We have targeted the first, and perhaps most essential skill that an empathic robot must possess: recognising emotions. The work presented here was based on the hypothesis that in order to develop machines that can express human-level intelligence, it is imperative to interpret existing knowledge of how the human body works and apply it to the computational models designed to provide robots with intelligence. Our work intends to reduce the gap between artificial and natural mechanisms by incorporating existing knowledge in the field of neuroscience into the development of artificial neural networks for emotion recognition and empathy imitation in a social robot. This will be achieved by making use of approaches such as features extraction, recognition and empathy that associates recognition and production in the manner found in the biological systems. In this paper, we addressed what is considered to be the first step towards empathic robots, emotion recognition.

Acknowledgements We would like to thank the School of Computing, Electronics and Mathematics, at Coventry University for funding this research. We would like to also acknowledge the contribution of the Barry Gidden fund which partially funded this work. The authors would also like to acknowledge

the invaluable contribution of Maria Charalambous, Kenny Ruiz, and Ibrahim Alamakky and Danielle Labhardt in creating this paper.

Conflict of Interest: The authors declare that they have no conflict of interest.

References

1. Ruiz-Garcia, A., Elshaw, M., Altahhan, A., and Palade, V. (2016) 'Emotion Recognition Using Facial Expression Images for a Robotic Companion'. in Engineering Applications of Neural Networks: 17th International Conference, EANN 2016, Aberdeen, UK, September 2-5, 2016, Proceedings [online] 79-93. available from <http://link.springer.com/10.1007/978-3-319-44188-7_6>
2. Lundqvist, D., Flykt, A., & Öhman, A. (1998). 'The Karolinska Directed Emotional Faces - KDEF'. CD ROM from Department of Clinical Neuroscience, Psychology section, Karolinska Institutet, ISBN 91-630-7164-9
3. Lucey, P., Cohn, J. F., Kanade, T., Saragih, J., Ambadar, Z., and Matthews, I. (2010). 'The Extended Cohn-Kanade Dataset (CK+): A complete expression dataset for action unit and emotion-specified expression'. Proceedings of the Third International Workshop on CVPR for Human Communicative Behavior Analysis (CVPR4HB 2010), San Francisco, USA, 94-101.
4. Ioffe, S. and Szegedy, C. (2015) 'Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift'. [online] available from <<http://arxiv.org/abs/1502.03167>>
5. Kanade, T., Cohn, J. F., and Tian, Y. (2000). 'Comprehensive database for facial expression analysis'. Proceedings of the Fourth IEEE International Conference on Automatic Face and Gesture Recognition (FG'00), Grenoble, France, 46-53.
6. Scassellati, B., Admoni, H. and Mataric, M. (2012) 'Robots for use in autism research'. Annual Review of Biomedical Engineering, 14, 275-294
7. Fasola, J. and Mataric, M. (2013) 'A socially assistive robot exercise coach for the elderly'. Journal of Human-Robot Interaction, 2.2, 3-32
8. Chang, W-L., Šabanovic, S. and Huber, L. (2013) 'Use of seal-like robot PARO in sensory group therapy for older adults with dementia'. Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction. IEEE Press.
9. Soler, M.V., Agüera-Ortiz, L., Rodríguez, J.O., Rebolledo, C.M., Muñoz, A.P., Prez, I.R., Ruiz, E.O., Sánchez, A.B., Cano, V.H., Chilln, L.C. and Ruiz, S.F., (2015) 'Social robots in advanced dementia'. Frontiers in aging neuroscience, 7
10. Robertson, J. (2014) 'Human Rights VS. Robot Rights: Forecasts from Japan'. Critical Asian Studies [online] 46 (4), 571-598. available from <<http://www.tandfonline.com/doi/abs/10.1080/14672715.2014.960707>>
11. Dahl, T. and Boulos, M. (2013) 'Robots in Health and Social Care: A Complementary Technology to Home Care and Telehealthcare?'. Robotics, 3, 1-21
12. Castellano, G., Paiva, A., Kappas, A., Nabais, F., Aylett, R., Barendregt, W., Hastie, H., and Bull, S. (2013) 'Towards Empathic Virtual and Robotic Tutors'. Lecture Notes in Computer Science. 733-736
13. Toh, L.P.E., Causo, A., Tzuo, P.W., Chen, I.M., Yeo, S.H., and Link, S. (2016) 'A Review on the Use of Robots in Education and Young Children'. Science Direct. 19, 148-163
14. Kory Westlund, J., Gordon, G., Spaulding, S., Lee, J., Plummer, L., Martinez, M., Das, M. and Breazeal, C. (2015) 'Learning a second language with a socially assistive robot'. Almere, The Netherlands
15. Affectiva (2016) [online] available from <<http://www.affectiva.com/>>
16. Kim, E.S., Berkovits, L.D., Bernier, E.P., Leyzberg, D., Shic, F., Paul, R., and Scassellati, B. (2013) 'Social Robots as Embedded Reinforcers of Social Behavior in Children with Autism'. Journal of Autism and Developmental Disorders 43 (5), 1038-1049
17. Rabbitt, S.M., Kazdin, A.E., and Hong, J.H. (2015) 'Acceptability of Robot-Assisted Therapy for Disruptive Behavior Problems in Children'. 101-110
18. KSERA. (2016) 'Knowledgeable Service Robots for Aging'. [online] available from <http://www.aat.tuwien.ac.at/ksera/index_en.html>
19. GeriJoy. (2016) 'Care and Companionship for Seniors - GeriJoy'. [online] available from <<http://www.gerijoy.com/>>
20. Leite, I., Pereira, A., Mascarenhas, S., Martinho, C., Prada, R. and Paiva, A., (2013) 'The influence of empathy in humanrobot relations'. in International journal of human-computer studies, 71(3), 250-260

21. Graaf, M.M.A., Ben Allouch, S. and Dijk, J.A.G.M., (2016) 'Long-term acceptance of social robots in domestic environments: Insights from a user's perspective'. AAAI
22. Duffy, B.R. (2006) 'Fundamental issues in social robotics'. *International Review of Information Ethics* 6
23. Boughrara, H., Chtourou, M., Ben Amar, C., and Chen, L. (2014) 'Facial expression recognition based on a MLP neural network using constructive training algorithm'. *Multimed Tools Appl.* 75, 709-731
24. Kahou, S., Michalski, V., Konda, K., Memisevic, R., and Pal, C. (2015) 'Recurrent Neural Networks for Emotion Recognition in Video'. *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction (ICMI '15)*. 467-474
25. Lawrence, S., Giles, C.L., Tsoi, A.C., and Back, A.D. (1997) 'Face Recognition: A Convolutional Neural-Network Approach'. *IEEE Transactions on Neural Networks* [online] 8 (1), 98-113. available from <<http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=554195>>
26. Gupta, A. and Garg, M. (2016) 'A human emotion recognition system using supervised self-organising maps'. *2014 International Conference on Computing for Sustainable Global Development (INDIACom)*. 654-659
27. Sarnarawickrame, K. and Mindya, S. (2013) 'Facial expression recognition using active shape models and support vector machines'. *Advances in ICT for Emerging Regions (ICTer)*, 2013 International Conference on. 51-55
28. Sohail, A.S.M. and Bhattacharya, P. (2011) 'Classifying Facial Expressions Using Level Set Method Based Lip Contour Detection and Multi-Class Support Vector Machines'. *International Journal of Pattern Recognition and Artificial Intelligence* [online] 25 (6), 835-862. available from <<http://www.worldscientific.com/doi/abs/10.1142/S0218001411008762>>
29. Ouellet, S. (2014) 'Real-Time Emotion Recognition for Gaming Using Deep Convolutional Network Features'. *CoRR* [online] abs/1408.3. 6. available from <<http://arxiv.org/abs/1408.3750>>
30. Ahsan, T., Jabid, T., and Chong, U.-P. (2013) 'Facial Expression Recognition Using Local Transitional Pattern on Gabor Filtered Facial Images'. *IETE Technical Review* [online] 30 (12), 47. available from <<http://tr.ietejournal.org/article.asp?issn=0256-4602;year=2013;volume=30;issue=1;spage=47;epage=52;aulast=Ahsan>> I slightly outperforms Gabor filters as a feature extraction method, at least for this particular dataset.
31. Mehta, N., and Jadhav, S. (2016) 'Facial Emotion recognition using Log Gabor filter and PCA'. *Proceedings of International Conference on Computing Communication Control and automation (IC-CUBE)*, 1-5.
32. Kanade, T., Cohn, J. F., and Tian, Y. (2000). 'Comprehensive database for facial expression analysis'. *Proceedings of the Fourth IEEE International Conference on Automatic Face and Gesture Recognition (FG'00)*, Grenoble, France, 46-53.
33. Chelali, F.Z. and Djeradi, A. (2015) 'Face Recognition Using MLP and RBF Neural Network with Gabor and Discrete Wavelet Transform Characterization: A Comparative Study'. *Mathematical Problems in Engineering* [online] 2015, 1-16. available from <<http://www.hindawi.com/journals/mpe/2015/523603/>>
34. Paul, P.P., Monwar, M.M., Gavrilova, M.L., and Wang, P.S.P. (2010) 'Rotation Invariant Multiview Face Detection Using Skin Color Regressive Model and Support Vector Regression'. *International Journal of Pattern Recognition and Artificial Intelligence* [online] 24 (8), 1261-1280. available from <<http://www.worldscientific.com/doi/abs/10.1142/S0218001410008391>>
35. Khan, S.A., Hussain, A., Usman, M., Nazir, M., Riaz, N., and Mirza, A.M. (2014) 'Robust Face Recognition Using Computationally Efficient Features'. *Journal of Intelligent & Fuzzy Systems* [online] 27 (6), 3131-3143. available from <<GotoISI>://WOS:000345981600037>
36. Hassner, T., Harel, S., Paz, E., and Enbar, R. (2015) 'Effective Face Frontalization in Unconstrained Images'. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* [online] 4295-4304. available from <<http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7299058>>
37. Hewahi, N.M. and Baraka, A.R.M. (2012) 'Impact of Ethnic Group on Human Emotion Recognition Using Backpropagation Neural Network'. *BRAIN. Broad Research in Artificial* [online] 20-27. available from <<http://brain.edusoft.ro/index.php/brain/article/view/284>>
38. Khashman, A. (2009) 'Application of an Emotional Neural Network to Facial Recognition'. *Neural Computing and Applications* 18 (4), 309-320
39. Burkert, P., Trier, F., Afzal, M.Z., Dengel, A., and Liwicki, M. (2015) 'DeXpression: Deep Convolutional Neural Network for Expression Recognition'. *arXiv preprint* [online] 1-8. available from <<http://arxiv.org/abs/1509.05371>>

40. Beaudry, O., Roy-Charland, A., Perron, M., Cormier, I., and Tapp, R. (2014) 'Featural Processing in Recognition of Emotional Facial Expressions'. *Cognition & emotion* [online] 28 (3), 416-32. available from <<http://www.ncbi.nlm.nih.gov/pubmed/24047413>>
41. Tan, C.C. and Eswaran, C. (2010) 'Reconstruction and Recognition of Face and Digit Images Using Autoencoders'. *Neural Computing and Applications* 19 (7), 1069-1079
42. Kuhn, H. and Tucker, A. (1951) 'Nonlinear Programming'. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. 481-492
43. Aldebaran. (2016) 'Who is NAO?'. [online] <<https://www.aldebaran.com/en/cool-robots/nao>>