Citation: Nikhil Churamani, Matthias Kerzel, Erik Strahl, Pablo Barros, Stefan Wermter.: Teaching Emotion Expressions to a Human Companion Robot using Deep Neural Architectures. International Joint Conference on Neural Networks (IJCNN), pages 627--634 - May 2017.

Teaching Emotion Expressions to a Human Companion Robot using Deep Neural Architectures

Nikhil Churamani, Matthias Kerzel, Erik Strahl, Pablo Barros and Stefan Wermter
Knowledge Technology Institute
Department of Informatics
University of Hamburg
Hamburg, Germany
Email:{5churama, kerzel, strahl, barros, wermter}@informatik.uni-hamburg.de

***Abstract*—Human companion robots need to be sociable and responsive towards emotions to better interact with the human environment they are expected to operate in. This paper is based on the Neuro-Inspired COmpanion robot (NICO) and investigates a hybrid, deep neural network model to teach the NICO to associate perceived emotions with expression representations using its on-board capabilities. The proposed model consists of a Convolutional Neural Network (CNN) and a Self-organising Map (SOM) to perceive the emotions expressed by a human user towards NICO and trains two parallel Multilayer Perceptron (MLP) networks to learn general as well as person-specific associations between perceived emotions and the robot's facial expressions.**

I. Introduction

As Human-Robot Interaction (HRI) research advances, artificial agents are expected to become an integral part of the human-centred environment. From performing daily duties to becoming true companions, agents need to perform a plethora of complex tasks while being socially responsive and even socially evocative. Thus, apart from possessing enhanced computing capabilities, agents also need to be able to factor in emotions and sentiments in their decision-making processes while operating in their human environments. It is thus quintessential that these agents are able to operate as well as interact intelligently with their surroundings.

Intelligence, in the context of an artificial agent, can be looked at as a measure of its ability to interact intuitively with humans and other agents [1]. The pivotal focus of much of HRI design is to make human-robot interactions as indistinguishable as possible from human-human interactions. As humans use many verbal as well as non-verbal cues such as facial gestures and body movement to communicate, the agent should not only be able to apprehend such inputs but also to express itself in a similar manner. It thus becomes important for agents to be able to model emotions and understand their interplay in order to enhance the interaction.

One of the major challenges in designing such an agent dealing with emotional feedback is the lack of standardisation in our understanding of emotions. Ranging from emotions being interpretations of physiological responses [2] to cognitive appraisal [3], many ideas have been proposed to define emotions over the years. Other studies [4] explain

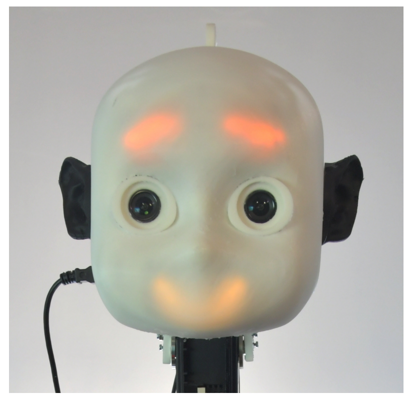

Fig. 1: NICO (Neuro-Inspired COmpanion), a robot platform for neuro-cognitive research.

emotions as evaluative judgements of the environment and other social agents with respect to the agent's goals and beliefs. Despite various contradictions, emotion modelling in agents can be proposed to be consisting of three phases contributing towards the Perception-Action cycle [5] for an emotional agent:

1) Emotion Perception: The agent observes the environment consisting of human and other artificial agents expressing a particular emotion. For example, a human smiling at the agent.
2) Emotion Synthesis: The agent factors in its own *goals and beliefs* to estimate an emotional state for itself. This is based on the inference engine of the agent so as to react to the perceived input from the environment.
3) Emotion Expression: Once the agent has received an input from the environment, it then expresses its emotional state in the form of facial gestures, synthetic speech etc. evoking another response from the environment.

The Perception-Action cycle [5] plays an important role in modelling instinctive human-robot interactions. It is extremely important for emotional agents to also be able to express emotions in order to be as natural as possible in an interaction and invoke further responses from the user. Agents like *Kismet* [6] take into account a variety of factors such as active behaviour, stimulus-based drive, affective appraisal etc. While performing affective interactions with humans, *Kismet* not only considers the affective state of the user but also the intensity and relevance of the stimulus it receives.

Expression can be grounded in robots based on the environment they operate in and their embodied capabilities. Approaches like the SAIBA [7] and SIRE [8] frameworks, aim to establish universal representations which can be customised to the robot platform being used. In the context of companion robots, they are always in the close vicinity of humans which acts as an important factor in deciding upon their design and operational capabilities. Our Neuro-Inspired COmpanion robot (NICO) (Fig. 1) acts as an example for a sociable robot that can be made to perform various day-to-day tasks and at the same time appear sociable to humans making use of its upper torso capabilities, synthetic speech or facial expressions. This study explores facial expressions of NICO (Fig. 1) as a way to encode emotion expression representations.

The proposed model evaluates the complete Perception-Action cycle for NICO as an emotional agent attempting to come up with expressions (more generally, action representations) that best relate to an emotional input. The model allows the robot to adapt to a user by customising its perception engine according to how the user expresses different emotions. It then attempts to associate the perceived emotion with an appropriate expression. This makes the agent both active and adaptive to the environment in which it operates. The model builds upon the perception model of Barros et al. [9] and attempts to train two parallel Multilayer Perceptron (MLP) network branches to associate different expression states to perceived emotions using a reward-based mechanism. The perception engine uses a combination of a Convolutional Neural Network (CNN) and a Self-organising Map (SOM) to recognise an emotion and learns to express the same using the MLP branches. The model explores Reinforcement Learning [10] to train NICO in an interactive and evaluative manner. Taking inspiration from the Training an Agent Manually via Evaluative Reinforcement (TAMER) algorithm [11], another solution was explored in our earlier work [12] which also aims at realising a "learning using human reward" scenario where the human drives the learning process by providing evaluative reward signals by mimicking the agent while interacting with it. This paper investigates a more refined scenario, and thus requires a more complex and robust solution than the one provided in our earlier work.

This paper first introduces NICO and explains its capabilities in Section II. It then details the proposed model in Section III, describing the different components of the hybrid neural network model. Section IV details the experiments conducted, both to define facial gestures for NICO as well as to train the model to associate these facial expressions with different emotions. Section IV also presents the results obtained in the experiments with a discussion on the same in Section V. Finally, Section VI concludes with the outcomes of the study and presents the future roadmap.

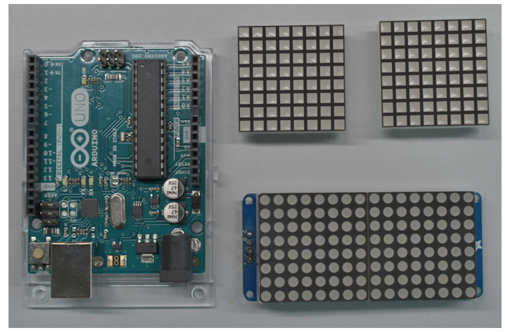

Fig. 2: NICO face expression hardware (Arduino microprocessor, two 8x8 matrix LEDs and a 16x8 matrix LED)

II. NICO - Neuro Inspired Companion

NICO, the Neuro-Inspired COmpanion (Fig. 1), is a one meter tall child-sized developmental robot designed and built for neuro-cognitive research. NICO features rich sensory, motor and facial expression capabilities which include, but are not limited to, bimanual manipulation, ability to communicate verbally, via hand and body postures, cameras for stereo vision, microphones embedded in pinnae for stereo hearing and haptic perception in the hands. It is designed and built to interact with humans as a companion in a natural environment where it can learn from experience and instruction. The flexible and modular design of NICO is adaptable to individual experimental set-ups, filling an existing gap for developmental humanoid robot platforms for embodied neuro-cognitive research.

A. *NICO Head Design*

The shape of the NICO head is based on the open-source design of the iCub surface [13]. This design aims to present pleasant, yet abstract and neutral features that avoid the uncanny valley effect. Onto this neutral surface of the NICO, emotional expression can be displayed using three LED arrays in the mouth and eye areas. These arrays enable NICO to show stylised facial expressions during human-robot interaction. Fig. 4 shows emotional expressions for the seven Universal Emotions (including *Neutral*) proposed by Ekman [14]: *Neutral, Happiness, Sadness, Anger, Surprise, Fear* and *Disgust*.

This facial expression system contains two 8x8 LED matrices for the eyebrows and an 8x16 LED matrix for the mouth (Fig. 2). The matrices are connected via an I2C bus with an Arduino micro-controller board [1] and can be controlled individually for each LED in the Matrix, which results in high flexibility for defining and modelling the expressions. The modelling of different expression representations on the face of NICO was inspired by recent fMRI studies which have shown that even text-based emoticons are based on the same response mechanism in the brain as natural faces [15]. Thus, the expressions were modelled on well-known text-based emoticons, but extended for more complex facial

[1] https://www.arduino.cc/

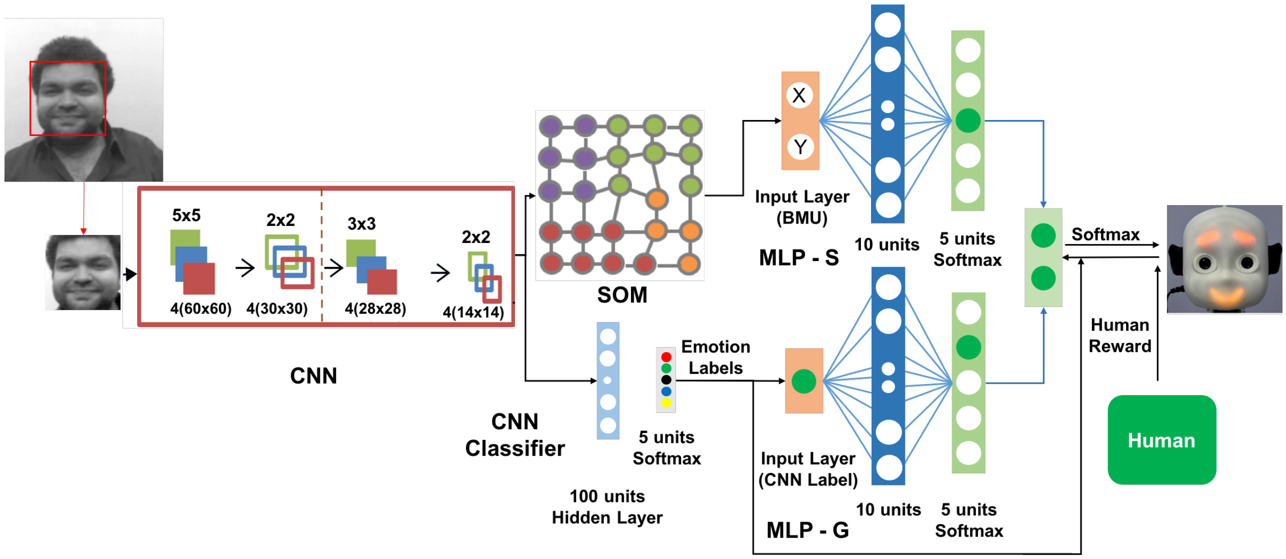

Fig. 3: Complete Model with CNN face representation, SOM clusters and MLP branches learning to associate an expression with the perceived emotion.

expressions such as *Disgust, Surprise* and *Fear*, which were modelled along the lines of the facial expressions on the iCub robot (which also uses LEDs) and the Flobi, which has a mechanical face expression modelling system [16].

The NICO head features visual sensors that offer access to vision similar to that of humans and thus offers a good platform to develop and evaluate bio-inspired perception models in multi-modal human-robot interaction scenarios. Two parallel mounted cameras with 2 Megapixel sensors offer a lens-adjusted field of view of 70°, shaped similarly to that of humans.

III. Proposed Model for Teaching Expressions

One of the most intuitive approaches towards learning to interact with the environment is to perform an action and evaluate how the environment reacts to that action [10]. The motivation for this study is to explore the possibility of teaching emotion expressions to NICO while interacting with the user. People express emotions differently and the robot should be able to adapt to this variance and customise its perception model so as to associate these emotions with appropriate action representations (expressions).

The model is based on the idea that a robot can be taught to express different emotions by simply by means of a human interacting with it and rewarding (or penalising) it for each action performed. The complete model (Fig. 3) is a deep, hybrid neural network consisting of three major components:

A. *CNN for Facial Feature Extraction*

The CNN layers of the model are used as a facial feature extractor. The agent uses them to form internal representations of human faces, each representing a particular emotion. The CNN is pre-trained on the Extended Cohn-Kanade [17] dataset and further trained on the data acquired during user studies with about 1000 images for each emotion so as to improve its performance. The network consists of two convolution layers (with max-pooling) trained using a dropout probability of 0.5, also making use of the L1 and L2 normalisation for each layer, respectively. The network parameters were inspired from the work of Barros et al. [9] and adapted empirically. The CNN is used to extract facial feature representations in the form of 784-dimensional feature vectors (the output of the second max-pooling layer). Each of the facial feature vectors encodes different emotion representations which are then used to train the SOM. The classifier from the CNN is used in parallel to classify emotional input and is used to provide ground-truth to the deeper layers of the network.

B. *SOM for Emotion clusters*

Once the facial feature representations are extracted using the CNN, they are used to train the SOM layer where it is expected that clusters emerge, each pertaining to a particular emotion. Whenever a human expresses an emotion towards the robot, the robot uses the CNN to represent the face as a feature vector and feeds the vector to the SOM as input. In the SOM, different clusters of neurons respond to different emotions and thus, the Best Matching Unit (BMU) provides the information to the deeper layers as to which emotion was expressed by the user. The model associates a facial feature vector with a BMU which acts as input for the MLP model. The proposed model uses a SOM lattice of 20×20 neurons (determined empirically) pre-trained on the Extended Cohn-Kanade [17] dataset and the training images acquired during user studies.

C. *MLP for learning Expression Representations*

The Multilayer Perceptron model consists of two branches, namely, MLP-G and MLP-S (Fig. 3). The generic MLP (MLP-G) uses the CNN classifier to provide ground-truth in terms of which emotion is being expressed by the user. It is trained using the CNN classifier labels, learning to associate these labels with appropriate action representations.

MLP-G can be thought of as a general action representation model associating expressions with a general perception of an emotion, indifferent of personal nuances. It adapts to how an emotion, in general, is expressed and learns to react to it. The specific MLP (MLP-S), on the other hand, uses the BMU from the SOM as input. It adapts to how a particular individual expresses an emotion. The SOM is able to represent this variance using the different clusters, each of which represents a particular emotion. Different BMUs fire for different individuals (sometimes even for the same individual depending upon how an emotion is expressed) for a given emotion and the model learns to associate an appropriate action with the BMUs. The MLP-S can be thought of being an action representation model associating expressions with a specific perception of an emotion as expressed by a particular individual, thus adapting to a particular user.

Both the MLP-G and MLP-S branches process the input in parallel and provide a probability value (using their individual softmax layers) to each of the expression representations. The output of each of these networks is again processed using another softmax operation resulting in the most probable expression which can be associated with the input emotion. The resulting expression is then shown to the user using the facial expression capabilities of the NICO. The ground-truth is provided by the CNN classifier generating the target label for each of the networks.

The human then evaluates this expression by either giving a positive, reaffirming reaction (expressing happiness) for a correct action, or a negative, reprimanding reaction (expressing anger) for an incorrect one. The given rewards are then accounted for by updating the model (using back-propagation) for the given interaction either 10 times for a positive reward (agent exploits the correct action) or 100 times for a negative reward (leading to the robot entering into an exploration phase again). Given that the action space is limited in this scenario, a higher importance is given to the negative reward as a conservative measure to avoid early exploitation.

Since both the network branches are trained in parallel, it is expected that the network switches between both branches to generate the output. It first learns to associate the general perception model with expressions and then fine-tunes itself to specific perception models, adapting to the user interacting with the agent.

IV. Experiments and Results

The study involved two different sets of experiments which needed to be conducted in order to assure a fair evaluation of the model. The first experiment involved studying how good the design of the facial expression in NICO was. This was done to make sure that the expression representations were apprehensible for the users and there was no confusion in identifying which emotion was being expressed by the robot. Once expression representations were successfully evaluated, the expressions with the least confusion possibility were chosen for the second experiment where the model was trained to associate these expressions with corresponding emotions expressed by the users. This section details each of these experiments and the results obtained. All participants were informed in detail about the objectives of the experiments and gave written informed consent for publishing the experiment results obtained using their data. After the experiment, participants were debriefed and questions arising from the study were answered.

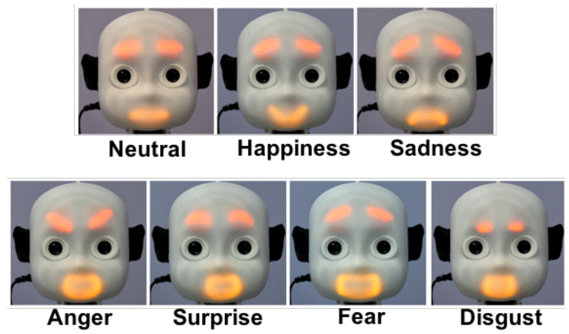

Fig. 4: NICO displaying facial expression for seven basic emotions.

A. Expression Representations for NICO

An empirical study with 20 participants (7 female, 13 male) was conducted to evaluate how well the seven universal emotions [14] (*Neutral, Happiness, Sadness, Anger, Surprise, Fear* and *Disgust*) could be conveyed with the NICO's facial emotion display (Fig. 4). All participants were aged between 19 and 49 years and reported English speaking abilities at least on a conversational level. The participants originated from eleven different countries from different regions of the world, including Europe, Asia, South America and the Middle East, providing a reasonable level of diversity to evaluate cultural nuances affecting emotion perception. As the experiment was aimed at evaluating the recognisability of the displayed emotions without any prior experience with the emotion display, four participants had to be excluded from the study due to their prior experience with the robot.

For the experiment, the seven emotions (encoded in facial expressions of NICO) were presented to the participants in a randomised order to avoid priming effects and to impede the establishment of a base line. Participants were instructed to look at the expression and document it as a multiple choice response and also to communicate their choice verbally. Participants were informed that they would be presented with seven different expressions but were not apprised as to the emotions they represent. Furthermore, participants were informed that they could select an emotion several times. They were also instructed to not correct their responses after having marked them once. Each emotion was displayed for 20 consecutive seconds. After 10 seconds, participants were reminded to mark their choice on the questionnaire and also verbally communicate the same. The whole experimental procedure took about 15 to 20 minutes for each participant.

1) Results: The responses from the participants (Fig. 5) show that a subset of 5 expressions (*Neutral, Happiness, Sadness, Surprise* and *Anger*) could be clearly identified by most (75% or above) of the participants. *Fear* was often

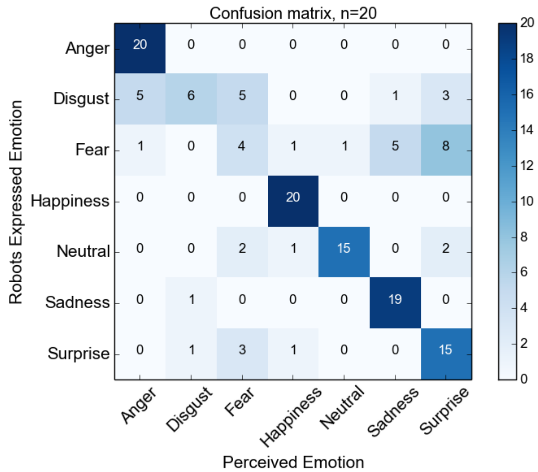

Fig. 5: Confusion matrix for human participants trying to recognise the emotion displayed by NICO.

confused with *Surprise*, and *Disgust* was poorly recognised, in general.

The results give evidence in favour of a subset of positive and negative emotions (*Anger, Happiness, and Sadness*) which were recognised by 95% of participants while *Neutral* and *Surprise* were recognised by 75% of them. This provides a reasonable foundation for the learning interactions described in the following sections, which use these 5 emotions to be learnt by the proposed learning model.

B. Teaching Expression Representations

Using the results obtained from the first experiment, the experiments for training the model were designed using 5 emotional states viz: *Anger, Happiness, Neutral, Sadness* and *Surprise*. The choice of *Anger, Happiness* and *Sadness* is quite intuitive as all the participants involved in the first experiment were able to correctly identify these emotions. Also, considering the Circumplex Dimensional Model [18] for emotion modelling, *Happiness* and *Sadness* mirror the valence of an emotion while *Anger* features on the high arousal side of negative valence. To have a more balanced evaluation, *Surprise* (high arousal and positive valence) and *Neutral* (no arousal and no valence) were also used in the evaluation, supported by the reasonably good identification rate from the first experiment. The expressions on NICO (LED combinations) were fixed based on the results of the first experiment. These could alternatively emerge as a result of a training mechanism as the robot tries out all possible combinations to express a particular emotion, leading to a huge search space. For the sake of simplicity and evaluation, this search space was limited to only 5 pre-defined LED combinations (subset of combinations from Fig. 4).

1) Experiment Setup: The experiment was conducted in two groups of 5 participants each. The participants from the first experiment were excluded to avoid any bias. The experiment conditions were kept the same for both the groups, as were all of the parameters of the learning algorithm such as *learning rate, number of interactions, network architecture*

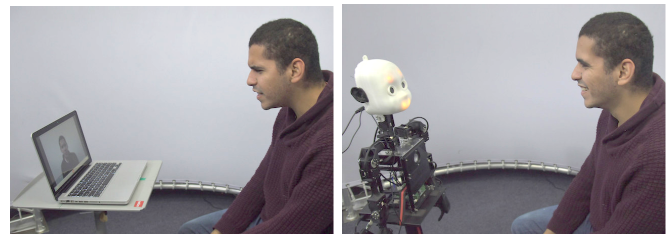

(a) Camera Booth for participants (b) Interaction with the Robot

Fig. 6: Experiment Preparation and Interaction

etc. The only variation for the different groups was that for the first group, the learning model (MLP-G and MLP-S) was always initialised randomly before the experiment and a new network was trained and evaluated for each participant. For the second group, the network was initialised once for the first participant and the training was continued for all the participants on the same model with evaluations conducted after each participant's interaction round was over.

Thus, for the first group the learning behaviour of the model while learning from scratch for different users was evaluated, while for the second group the effect of a continuous learning with different subjects was examined.

2) Preparation: Once the experiment setup was decided upon and the groups were formed, the procedure to be followed for the experiment with each of the participants was standardised. The experiment involved the participant expressing 5 emotions towards NICO and thus it was essential that the participants were comfortable in expressing these emotions as effortlessly as possible. The participants were advised to avoid any kind of exaggeration or 'acting' of emotions and to express the emotions as naturally as possible. Participants were provided with a camera booth (Fig. 6a) showing them a camera-feed of themselves to practice expressing the 5 emotions in case they felt it was needed before performing the experiment. Once the participants were ready to start the experiment, they were shown the five expression representations in NICO (resulting from the first experiment) and were made aware of the emotion each one represented.

3) Interaction Scenario: Each participant (in each group) was asked to perform 10 interaction rounds with the robot. In each interaction round, the participant was asked to express the 5 emotions (Fig. 6b). For each emotion, the interaction was split into two phases:

- Input Phase: In the input phase, the participant was asked to express a particular emotion (one of the 5 emotions viz. *Anger, Happiness, Neutral, Sadness* and *Surprise*) towards the robot. The robot captures an image of the participant and tries to perceive the emotion expressed using the CNN and SOM layers. It then uses this information to come up with an action representation using the MLP model (MLP-G or MLP-S) which best mimics the participant's emotion and expresses it using the facial LED combination defined for that emotion.
- Reward Phase: Once the robot comes up with a reaction

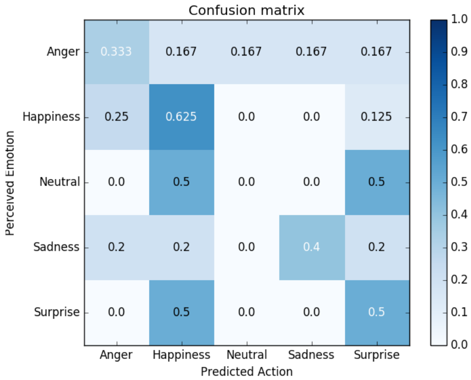

(a) MLP-S

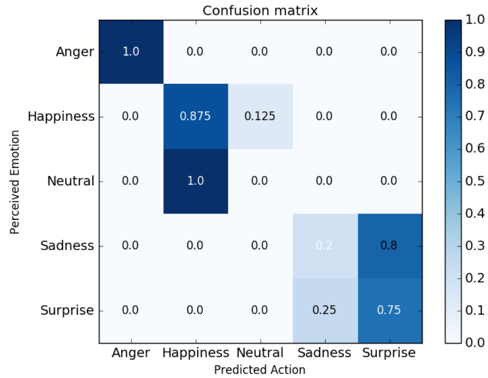

(b) MLP-G

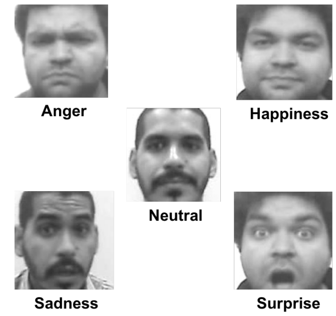

(c) Input Sample

Fig. 7: Sample network input and Normalised Confusion Matrices for the last epoch scores for 5 participants

to mimic the emotion expressed by the participant, it is then rewarded or penalised based on how accurate this imitation was. If the participant felt that the robot's expression was correct, the participant was asked to reward it using an affirmative expression (for example, *Happiness*) and if it was incorrect, the participant was asked to reprimand it using a negative expression (for example, *Anger*). This reward provided by the participant acts as a factor for deciding whether the robot should *exploit* a particular action space or it should *explore* for a better expression representation.

The above two phases were repeated for 5 emotions per interaction. Each time the participant was told which emotion to express in order to make sure all five emotions are included in each interaction round. The interaction scenario (for 5 emotions) was then repeated a total of 10 times to generate 50 interaction data samples to train the model. Data augmentation was implemented by recording the input and target labels for all the interactions and using them iteratively to update the model at the end of each interaction so as to speed up the learning process.

4) Results: As described earlier, the experiments were conducted with two user groups (with five participants in each group) keeping the same experiment setup with the only difference being how the network was trained for different participants within a user group. For the first group, for each participant, the MLP-G and MLP-S networks were trained from scratch initialising the weights randomly (drawing from a uniform distribution). For the second group, the MLP layers were randomly initialised only for the first participant and continued learning with each participant providing a large number of user interactions for training.

1) User Group 1: The objective of the experiment with the first user group was to investigate how the model learns with individuals and how different users drive the training. Also, another factor for investigation was the effect of the expressiveness of the participant on the training i.e. if the learning depended on how well an individual was able to express an emotion or if the model was able to adapt quickly to the participant's expressions. The resulting normalised confusion matrices for the last interaction for MLP-S (Fig. 7a) and MLP-G (Fig. 7b) for all the participants clearly show that the model, on average, was able to correctly associate a subset of emotions with expressions, but not all of them. Sample input images can be seen in Fig. 7c. Some participants were highly expressive, and in these cases the model was able to learn at least two emotions in the limited number of interactions. For other participants, the model was only able to learn one or at best two emotions, using both MLP-S and MLP-G. Nonetheless, the results show that the model, even with only ten interaction rounds, was able to associate expression representations with emotions. This is promising, as it hints towards continued learning i.e. continuing to train the same network for all participants, offering improved results and potentially learning more expression representations. Thus, we explored this factor in the experiment conducted with the second group.
2) User Group 2: For the second group of participants, the model was initialised randomly for the first participant who trained it for 10 interaction rounds. The following participant continued training the same network. There was an evaluation round for each of the participants between the transitions to see if and how the network learnt to express emotions. It was observed that the network showed continuous improvement in learning expressions, which is reflected in the F1-scores for each evaluation (Fig. 8) where an overall increase in these values can be seen as training continues for the model. The model showed an increasing growth in performance starting with learning 2 expression representations for the first participant to 4 expressions for the final participant. Another point to be noted here is the individual performance of the MLP-G and MLP-S branches. The MLP-G branch, the general perception model, is able to learn expressions very early and helps drive the training further (Fig. 8). The MLP-S soon catches up and is able to correctly predict most of the expressions by the fourth participant. The local minimum in the plot after the fourth participant

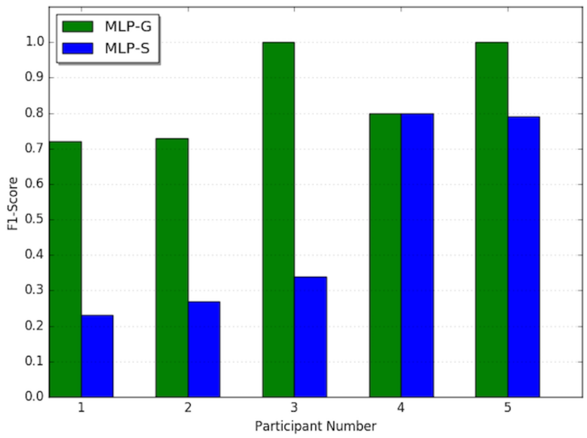

Fig. 8: F1-Score for continued training with 5 participants

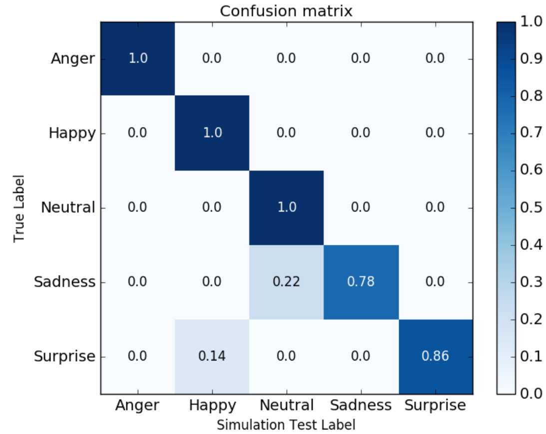

Fig. 9: Final Confusion Matrix after extending the learning by the Simulated Test for 100 additional interaction rounds

might indicate the network generalising for multiple participants but the overall positive trend suggests improvements in the learning.

3) Simulated Test: It was expected that continued learning from User Group 2, would result in the model learning to correctly express all the recognised emotions. Thus, we tried to run a simulated training where the network continued the training for another 100 interaction rounds using labelled images of participants expressing different emotions. For each interaction round, 5 images were chosen randomly and given as an input to the model with the labels coming from the true labels of the images used as the reward signal. It was observed that the network was able to learn all 5 emotions with high accuracy and precision (Fig. 9) after only 100 more interactions.

V. Discussion

The experiments offer promising results validating this study's objective of evaluating mechanisms for teaching a companion robot to express different emotions. The study based itself on evaluating how humans perceive facial emotion expressions (Section IV-A) and then, using these results, modelled emotion expression representations in the NICO. The proposed model was able to learn to associate these expression representations with the emotions expressed by the participants (Section IV-B) learning to adapt itself to individual nuances when it comes to facial expressions. The results from the continued learning approach (Fig. 8 and Fig. 9) provide evidence in favour of the model learning to perform well for different individuals with the general as well as the specific perception models adapting to the variance with which different participants expressed emotions. The general perception model (MLP-G) learns to associate general emotion perception (using the CNN labels) to action representations. The use of the word 'general' here is justified as the CNN classifier classifies all the face representations into five emotion classes ignoring the intra-class variance. This can be compared to perception of emotions in humans who can almost instantly differentiate between different emotions a particular individual expresses but who might find it complex to identify whether different individuals are expressing the same emotion. The specific perception model tries to address this issue by using the SOM lattice for clustering different facial features expressing different emotions. Each of these clusters encodes how different individuals express the same emotion and how different BMUs fire for different individuals expressing the same emotion. The proposed approach is different from ensemble methods as the network, rather than generalising using a number of classifiers, adapts to the variance with which different users express emotions, and learns to express the desired emotion.

For both the experiment groups (Group 1 and Group 2 including the Simulated Test) there were some emotions which were misclassified by the CNN which could be due to the variance in the expressiveness of different participants which was also explored in the study. The images of all the participants interacting with the model were recorded and 20 interaction rounds were randomly selected to be labelled both by the CNN and 10 independent human annotators. This was done to compare the classification results of the CNN with those of humans (Fig. 10 and Fig. 11). If a particular emotion is misclassified by a human observer, it could be argued that the CNN also learns in a similar way and is bound to make some classification errors as a result.

The results support our assumptions with the resultant Cohen's Weighted Kappa [19], [20] and Kendall Tau distance [21] values for CNN classification and the ground-truth values being $\kappa = 0.81$ and $K = 0.829$ respectively, against the mean Kappa and Tau distance values for the 10 annotators and the ground-truth values being $\kappa = 0.86$ and $K = 0.834$, respectively. Similar values indicate that human performance and the network performance are comparable when it comes to recognising the expressed emotions for different individuals.

VI. Conclusion and Future Work

Emotion perception and expression are desired capabilities in human companion robots. Companion robots need to be emotionally responsive in order to make interactions more

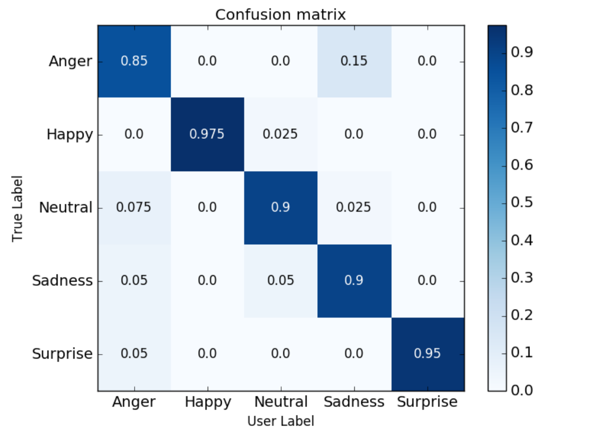

Fig. 10: Normalised Confusion Matrix for Annotator Labels

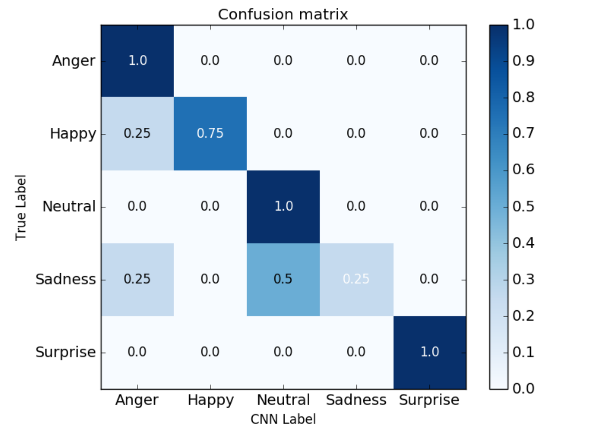

Fig. 11: Normalised Confusion Matrix for CNN Labels

intuitive. They should be able to pick up non-verbal cues to improve their decision making capabilities, and emotion modelling offers an interesting approach in this direction. The paper presents a model which can be trained on-the-fly to adapt to the human environment and to learn to recognise and express different emotions. This resonates with an important goal of HRI research of being able to train robots using non-expert users in an online fashion. As future work, it would be interesting to explore reinforcement learning models [22]–[24], in the realm of reward and policy shaping mechanisms and compare the results with the presented model. This model offers promising results and in the future can also be extended to a multi-sensory input scenario where both perception and expression models can be extended to include intoned speech, body language or even haptic feedback.

ACKNOWLEDGEMENT

This work was partially supported by the German Research Foundation (DFG) under project Cross-modal Learning (TRR 169) and the Hamburg Landesforschungsförderungsprojekt.

REFERENCES

[1] A. M. Turing, "Computing machinery and intelligence," *Mind*, pp. 433–460, 1950.
[2] W. B. Cannon, "The James-Lange theory of emotions: A critical examination and an alternative theory," *The American Journal of Psychology*, pp. 106–124, 1927.
[3] R. S. Lazarus, J. R. Averill, and E. M. Opton, "Towards a cognitive theory of emotion," *Feelings and Emotions*, pp. 207–232, 1970.
[4] E. Hudlicka, "What are we modeling when we model emotion?" in *AAAI Spring Symposium: Emotion, Personality, and Social behavior*, 2008, pp. 52–59.
[5] J. M. Fuster, "Upper processing stages of the perception–action cycle," *Trends in cognitive sciences*, vol. 8, no. 4, pp. 143–145, 2004.
[6] C. Breazeal, "Emotion and sociable humanoid robots," *International Journal of Human-Computer Studies*, vol. 59, no. 1, pp. 119–155, 2003.
[7] S. Kopp, B. Krenn, S. Marsella, A. N. Marshall, C. Pelachaud, H. Pirker, K. R. Thórisson, and H. Vilhjálmsson, "Towards a common framework for multimodal generation: The behavior markup language," in *Intelligent Virtual Agents. IVA 2006. Lecture Notes in Computer Science, vol 4133.*, J. Gratch, M. Young, R. Aylett, D. Ballin and P. Olivier, Ed., 2006.
[8] A. Lim and H. G. Okuno, "The mei robot: towards using motherese to develop multimodal emotional intelligence," *Autonomous Mental Development, IEEE Transactions on*, vol. 6, no. 2, pp. 126–138, 2014.
[9] P. Barros and S. Wermter, "Developing crossmodal expression recognition based on a deep neural model," *Adaptive Behavior*, vol. 24, no. 5, pp. 373–396, 2016.
[10] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT Press, 1998.
[11] W. B. Knox and P. Stone, "TAMER: Training an agent manually via evaluative reinforcement," in *IEEE 7th International Conference on Development and Learning*. IEEE, 2008, pp. 292–297.
[12] N. Churamani, F. Cruz, S. Griffiths, and P. Barros, "iCub: Learning Emotion Expressions using Human Reward," in *Workshop on Bio-inspired Social Robot Learning in Home Scenarios, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Daejeon, Korea, 2016.
[13] G. Metta, L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. von Hofsten, K. Rosander, M. Lopes, J. Santos-Victor, A. Bernardino, and L. Montesano, "The iCub Humanoid Robot: An Open-Systems Platform for Research in Cognitive Development," *Neural Networks*, vol. 23, no. 8–9, pp. 1125–1134, 2010.
[14] P. Ekman, "An argument for basic emotions," *Cognition & Emotion*, vol. 6, no. 3-4, pp. 169–200, 1992.
[15] K. W. Kim, S. W. Lee, J. Choi, T. M. Kim, and B. Jeong, "Neural correlates of text-based emoticons: a preliminary fMRI study," *Brain and Behavior*, vol. 6, no. 8, 2016.
[16] F. Hegel, F. Eyssel, and B. Wrede, "The social robot 'Flobi': Key concepts of industrial design," in *19th International Symposium in Robot and Human Interactive Communication*. IEEE, 2010, pp. 107–112.
[17] P. Lucey, J. F. Cohn, T. Kanade, J. Saragih, Z. Ambadar, and I. Matthews, "The Extended Cohn-Kanade Dataset (CK+): A complete expression dataset for action unit and emotion-specified expression," in *Proceedings of the Third International Workshop on CVPR for Human Communicative Behavior Analysis (CVPR4HB 2010)*, San Francisco, USA, 2010, pp. 94–101.
[18] J. A. Russell, "A circumplex model of affect," *Journal of Personality and Social Psychology*, vol. 39, pp. 1161–1178, 1980.
[19] J. Cohen, "Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit," *Psychological Bulletin*, vol. 70, no. 4, p. 213, 1968.
[20] J. Carletta, "Assessing agreement on classification tasks: the kappa statistic," *Computational linguistics*, vol. 22, no. 2, pp. 249–254, 1996.
[21] M. G. Kendall, "Rank Correlation Methods," 1948.
[22] A. L. Thomaz and C. Breazeal, "Teachable robots: Understanding human teaching behavior to build more effective robot learners," *Artificial Intelligence*, vol. 172, no. 6, pp. 716–737, 2008.
[23] M. Grzes and D. Kudenko, "Reinforcement learning with reward shaping and mixed resolution function approximation," *International Journal of Agent Technologies and Systems (IJATS)*, vol. 1, no. 2, pp. 36–54, 2009.
[24] S. Griffith, K. Subramanian, J. Scholz, C. Isbell, and A. L. Thomaz, "Policy shaping: Integrating human feedback with reinforcement learning," in *Advances in Neural Information Processing Systems 26*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2013, pp. 2625–2633.