

“Why the face?”: Exploring Robot Error Detection Using Instrumented Bystander Reactions

Maria Teresa Parreira

mb2554@cornell.edu

Cornell University, Cornell Tech
New York, USA

Ruidong Zhang

Cornell University
New York, USA

Sukruth Gowdru Lingaraju

Cornell Tech
New York, USA

Alexandra Bremers

Cornell Tech

New York, USA

Xuanyu Fang

Cornell Tech

New York, USA

Adolfo Ramirez-Aristizabal

Accenture Labs
SF, USA

Manaswi Saha

Accenture Labs

SF, USA

Michael Kuniavsky

Accenture Labs

SF, USA

Cheng Zhang

Cornell University

New York, USA

Wendy Ju

Cornell University, Cornell Tech

New York, USA

ABSTRACT

How do humans recognize and rectify social missteps? We achieve social competence by looking around at our peers, decoding subtle cues from bystanders — a raised eyebrow, a laugh — to evaluate the environment and our actions. Robots, however, struggle to perceive and make use of these nuanced reactions. By employing a novel neck-mounted device that records facial expressions from the chin region, we explore the potential of previously untapped data to capture and interpret human responses to robot error. First, we develop NeckNet-18, a 3D facial reconstruction model to map the reactions captured through the chin camera onto facial points and head motion. We then use these facial responses to develop a robot error detection model which outperforms standard methodologies such as using OpenFace or video data, generalizing well especially for within-participant data. Through this work, we argue for expanding human-in-the-loop robot sensing, fostering more seamless integration of robots into diverse human environments, pushing the boundaries of social cue detection and opening new avenues for adaptable robotics.

KEYWORDS

wearable; robot error; social signals; HRI; computer vision; human-AI collaboration

1 INTRODUCTION

Humans navigate social environments by perceiving and interpreting cues from those around them. *Social competence* stems partly from this ability to read non-verbal signals such as facial expressions, body language, and subtle gestures. These environmental and peer-based inputs form a feedback loop that informs social behavior and facilitates integration into diverse social contexts [18].

The concept of *human-in-the-loop sensing* in robotics leverages this innate human capability to enhance the social competence and




Figure 1: User study scheme. Participants wearing NeckFace [8] watch videos where a scenario of human or robot error is shown, eliciting a reaction. The IR camera image is converted into 3D facial points and head rotation data through a customized NeckNet model. This data is then used to train error detection models which map human reactions to the scenario displayed.

efficiency of robotic systems. This approach posits that human reactions can serve as a rich, real-time source of information for robots, guiding their behavior and functionality. For example, methods for Reinforcement Learning from Human Feedback (RLHF) have leveraged human social cues and emotions as feedback input [24]. Recent works [31, 36, 39] have started exploring the potential of human reactions in the context of human-robot interaction (HRI) for robot error detection. These studies make use of human social cues, captured through various means, as indicators of robot performance and social appropriateness.

Traditional implementations of this concept often rely on stationary cameras or extensive sensor arrays to capture human reactions. While groundbreaking, these methods can be resource-intensive, computationally demanding, and may struggle to adapt to dynamic, real-world environments where both humans and robots are in constant motion. To address these limitations, we employ NeckFace [8], a novel, neck-mounted device that records facial expressions from the chin region and maps them onto 3D facial points and head rotation. This lightweight, bi-camera system offers an alternative perspective on human reactions, potentially capturing nuanced expressions that may be overlooked by traditional methods. Its wearable nature allows for data collection in dynamic, real-world scenarios, providing a more efficient and adaptable solution for understanding human responses to robot behavior.

To test and benchmark the efficiency of NeckFace against other conventional approaches, we conducted a comparative user study. Participants' reactions to video stimuli displaying human and robot error were recorded simultaneously using NeckFace and conventional RGB cameras. Following this stage, we 1) developed and tested a model to map NeckFace's IR camera data onto 3D facial expressions and head motion data, NeckNet-18; 2) used this reaction data to develop error detection models that perform well for unseen participants but also generalize to a single participant with minimal data for training; and 3) compared these models with conventionally used methods, such as OpenFace features [3, 36, 39] or RGB camera data [4, 30], finding that it outperforms these methods.

This research not only contributes to the development of more dynamic and context-aware robotic systems but also advances our understanding of how human social competence can be effectively leveraged in HRI. By exploring alternative data sources and expanding on human-in-the-loop sensing methodologies, we aim to pave the way for more socially competent and efficient robotic systems that can seamlessly integrate into human environments.

2 RELATED WORK

Social cues for error detection: Human-in-the-loop robotic systems [35] leverage human input to enhance robots' efficiency, yielding positive outcomes in HRI [11]. Following this concept, harnessing human social cues can equip robots with better social and operational competence. For example, a substantial body of literature has begun exploring social cues in response to robot error, revealing that human reactions are multimodal and diverse [7, 31], including body and head motion [12, 20, 40], gaze [6], and facial expressions [1, 16, 20, 38]. The detection of robot error from social cues has been investigated through various machine learning approaches. Stiber et al. [39] proposed a framework for "error aware" HRI, employing a deep neural network, building on previous work [37]. Other studies [31, 36] utilize recurrent neural networks (RNNs) to preserve time-dependencies. Random Forest classifiers have also demonstrated efficiency as linear methods for error detection [26]. Recently, combinations of Minirocket [10] classifiers have achieved high performance in error and confusion detection [41, 42]. To bypass the need for feature extraction, some researchers have employed convolutional neural networks [4, 30] to classify human reactions directly from video data. Recent work by Ravishankar

et al. [33] explores unsupervised learning of robot error by detecting abnormalities in human-robot interactions from human-human interaction data. These works explore visual, audio and semantic data modalities. Given the novelty of the field, however, there remains untapped potential for exploring novel data modalities and processing methodologies for social cue classification.

Beyond standard methods for data collection: Capturing facial expressions as humans watch robots can allow the robot to infer human perceptions of robot performance and social appropriateness. High-accuracy facial expression tracking traditionally involves active motion capture instrumentation on the face, either using visual markers [34] or other sensors such as electromyography (EMG) [13] or capacitive sensors [32]. These methods usually require extensive setups and calibration and can cause physical and social discomfort. Relieving users from heavy on-face instrumentation, frontal-camera-based vision solutions have been getting recent attention [2, 27]. These solutions usually involve placing a camera in front of the user's face to directly record and analyze their facial expressions. While they provide a less intrusive experience, they are not mobile, as users have to appear in front of the camera within reasonable angle and distance. To address these challenges, recent wearable solutions look more into less intrusive sensing solutions while trying to maintain tracking performance [8, 22, 23, 43].

Among these, NeckFace [8] is a wearable facial expression tracking system that can estimate 3-dimensional facial expressions continuously. It performs comparably to traditional video-based systems, without the need for a frontal camera. Unlike other mounted devices, NeckFace is robust to users' head rotations, allowing for greater flexibility in use. It features a 3D-printed neckband with two infrared (IR) cameras positioned on either ends pointing upwards toward the wearer's neck, chin and cheeks (Figure 1). Two 850 nm IR LEDs are placed beside each camera to provide independent lighting to minimize the impact of unstable environmental lighting. NeckFace uses a convolutional neural network (ResNet-34 [15]) to convert the captured images into 3D facial expressions represented by 52 Blendshapes. NeckFace has been demonstrated to have reliable tracking performance across different scenarios such as sitting, walking and remounting. In light of this, we chose NeckFace as our facial reaction reconstruction device.

To the best of our knowledge, this is the first user study pursuing the use of a neck-mounted device to capture human social cues in HRI, investigating and benchmarking its potential for error detection.

3 STUDY DESIGN

In this section, we detail the NeckFace system implementation, user study design and methodology for data collection. To collect human reactions to observed errors, we followed a procedure adapted from Bremers et al. [4] for in-lab data collection, where participants observed a set of videos where error occurs. The elicitation materials and code used are shared in the project repository ¹.

3.1 Study Protocol

The study set up is shown in Figure 2. Participants were recruited through flyers and word-of-mouth. After giving informed consent

¹<http://irl.tech.cornell.edu/badrobots-feat-neckface/>



Figure 2: Study setup. In the *Calibration round*, the participant, wearing NeckFace [8], copies the movements seen on a video on an iPhone 11. The *Stimulus round* consists of a series of 30 videos played on a screen, while a webcam and NeckFace record facial reactions.

and providing demographic information, participants were taken to the experiment room, where the researcher mounted NeckFace on them, adjusting the positioning of the cameras as necessary. Following this, participants completed a *calibration round*, where they were placed in front of an iPhone device and asked to copy the movements seen on a video, such as moving their head up and down or raising their eyebrows (approximately 5 minutes), complying with the initial NeckFace work [8]. After this, the researcher would leave the room and data collection would begin, as participants completed a *stimulus round* where they watch a series of 30 videos presented in random order on a screen. Participants watched each stimulus video while a high-resolution webcam and NeckFace recorded their facial responses. Participants were not able to see their own image while the stimulus videos played. Compensation was provided at rate of USD 15/hour. After this, participants were taken back to a smaller room to fill out a final written questionnaire. The full procedure took around 30 minutes to complete. A diagram of study stages and data collected can be seen on Figure 3. This data was collected under Cornell University IRB exempt protocol #1609006604.

3.2 System implementation

In this work, both IR cameras from NeckFace are connected to a single Raspberry Pi 4B via an Arducam multi-camera adapter, different from the original NeckFace system which uses two Raspberry Pis to record the images. This allows for a lighter system that is easier to wear and integrate with other components. The two cameras capture images of the wearer’s neck, chin and cheeks from both sides, which are subsequently used to estimate the full facial expressions in 3D.

The NeckFace system needs to be trained in order to provide accurate facial expression estimations. To supervise the training process, a dataset of (IR cameras image, 3D facial expression) pairs needs to be collected. Specifically, the 3D facial expressions (ground truth) were represented by 52 Blendshapes [34] captured with the

TrueDepth camera on an iPhone 11 Pro. We also predicted three head rotation variables. We obtained the training data through the calibration round, where participants were asked to follow movements and facial expressions seen on a video (Figure 2).

In addition to the NeckFace device and calibration setup, a participant-facing RGB camera was used to record participants’ reactions.

3.3 Stimulus Dataset

We selected a set of 30 stimulus videos for collecting human reactions. Following the methodology in Bremers et al. [5], we define errors as “actions not as intended”. For example, a robot falling while playing soccer, or a person crashing a lawnmower onto an object. We included videos of human failures (10), robot failures (10) and control videos (10) (see Figure 2 for example). The videos’ average length is 13.69 ± 7.77 s. The full list can be consulted in the project repository¹.

3.4 Participants

A total of 30 participants completed the study. Due to technical issues or poor quality of data, 5 participants were excluded (P17 and P27 were removed because they were wearing large earrings which would shake and cause disturbances in the image quality; P27, 29 and 30 for inadequate NeckFace camera adjustment, and P13 due to erratic head motion during calibration). The remainder 25 participants took on average $33m22s \pm 4m19s$ to complete the survey. Ages range from 21-78 (37.73 ± 17.37). 17 participants identify as female, 12 as male, and 1 as non-binary. The study included participants from 11 nationalities. Racial/ethnical distribution includes 17 Caucasian/White or Asian/White, 6 Asian/Asian American participants, 4 Hispanic/Latino and 1 African/African American/Black, 2 participants who self-described to be Indian and Indian American.

4 MODEL DEVELOPMENT

Below, we describe model development for NeckFace 3D facial reconstruction and for error detection models.

4.1 3D Facial Mapping Model

The 3D facial mapping model takes IR video frames captured by the NeckFace device and converts them into numerical representations of 3D facial expressions. The model was developed following the original design of NeckNet [8].

Problem Formulation and State space: Tracking facial expressions with the wearable neckband was formulated as a regression problem: at any given time, the raw data from the wearable device is captured as two images $i_l, i_r \in \mathbb{R}_{640 \times 480}$, from the left and right side camera, respectively. The model output is a feature vector $e \in \mathbb{R}_{55 \times 1}$, representing the 3D facial expression and head orientation.

Action Space : The model output $e \in \mathbb{R}_{55 \times 1}$ represents the 3D facial expression. This includes 52 Blendshape parameters extracted from TrueDepth camera by ARKit, representing the frontal face, and 3 angles representing the orientation (yaw, pitch, roll) of the head. The 52 Blendshapes each represent a specific facial movement (e.g., jawOpen, mouthClose, etc.), with numerical values linearly




Figure 3: Study protocol and data collected. In the calibration round, NeckFace IR camera data is collected along with the TrueDepth data from the iPhone (and head rotation angles). The latter serves as ground truth to train NeckNet-18. In the stimulus round, reactions to neutral (0) and error (1) videos are collected through NeckFace cameras. This dataset, NeckFaceIR, is later input to NeckNet-18, transforming the dataset into 3D facial reactions (NeckData).

describing the extent of that movement from a scale of 0 to 1000. The ground truth of these Blendshapes was captured with the TrueDepth Camera of an iPhone 11 at 30 fps and later resampled to 12 fps to match the sampling rate of the IR cameras. A full list of the 55 features is provided as supplementary material.

Dataset : Two datasets were used to train and evaluate the facial reconstruction model. In addition to the data collected in this study, we used the original NeckFace dataset [8] for training.

- **NeckFaceOriginal**: video frames collected for Chen et al. [8]. It includes images recorded from two IR cameras at 30 fps from 13 participants. This dataset was used to pre-train the model.
- **NeckFaceIRCalib**: video frames from NeckFace’s IR cameras collected during the calibration round. Frames were recorded at 12 fps. This dataset was used to fine-tune the pre-trained model. See Figure 3 for schematic of how this data was used.

Data processing and augmentation were applied both during pre-training and fine-tuning. Specifically, raw images from the two cameras were first converted to grayscale and resampled to a size of 320 by 240 pixels, then tilted horizontally to generate an image of size 640x240 pixels. Data augmentation methods include: 1) 50% chance of random scaling with a factor of 0.9 to 1.1, 2) 50% chance of random rotation of -30 to +30 degrees (while pre-training) or -8 to +8 degrees (while fine-tuning), 3) 50% chance of random translation of -6 to +6 pixels in both directions. These operations were applied to both sides of the images independently but with the same parameters.

Evaluation Metrics : The model was evaluated with the Mean Absolute Error (MAE) between the predicted Blendshapes and ground truth (TrueDepth). The error of the facial expressions and

head orientation were calculated separately. Specifically, $MAE_f = \sum_{i=1}^{52} \frac{1}{52} |e_i - \hat{e}_i|$, $MAE_o = \frac{1}{3} \sum_{i=53}^{55} |e_i - \hat{e}_i|$ where MAE_f is the error of facial expressions tracking and MAE_o is the error of head orientation tracking. e is the ground truth of facial expression and head orientation and \hat{e} is the model prediction.

Model Design : We designed NeckNet-18, a convolutional neural network (CNN) with a ResNet-18 backbone and a fully-connected decoder (Figure 4). We based this structure on the previous work [8], with adaptations to better fit our data and task. First, we used a lighter-weight ResNet-18 instead of the ResNet-34 model used in prior work [8], based on two considerations: 1) the dataset collected in this user study is smaller than that in Chen et al. [8] due to smaller calibration time and a lower frame rate. Larger models would be more prone to overfitting. 2) ResNet-18 has roughly half the number of trainable parameters when compared with ResNet-34, which results in faster training and inference speed. Previous work on similar tasks also demonstrated that ResNet-18 has comparable performance to ResNet-34 with a smaller dataset [23].

Model Training : NeckNet-18 is designed for within-participant generalization, using a small Calibration round to fine-tune the model for new participants. Thus, the model’s generalizability relates to unseen *temporal* data rather than unseen participants.

We pre-trained the model on the *NeckFaceOriginal* dataset. To fine-tune and evaluate model performance of the 3D facial mapping model, we used data collected from the calibration round, *NeckFaceIRCalib*. Specifically, we first divided this dataset into 5 folds, by dividing data collected from each participant’s calibration round into 5 parts in temporal order, and combining each of the 5 parts from different participants together. This means that each fold contains data from all participants, but this data does not overlap temporally. We then chose 4 of the 5 folds to use as training data to fine-tune the model, and the other fold as testing data. This process was repeated 5 times, constituting a 5-fold cross-validation where all 5 folds had appeared in the testing data once. This step also allowed us to understand the within-participant generalizability of the model. Finally, all 5 folds were combined and the model was trained again to maximize the benefit of all available training data.

During pre-training, the model was trained for 30 epochs with an initial learning rate of 0.0002. Given the differences in camera angle and image quality between the original NeckFace implementation and our re-implementation, the fine-tuning step which followed was given more weight to allow the model to better learn the distribution of the new dataset. As a result, during fine-tuning, the model was also tuned for 30 epochs but with an initial learning of 0.0001.

After training, the model was used to reconstruct the NeckFace IR data collected during the *stimulus round* (participant reacted to video scenarios), which generated the dataset of 3D facial reactions used as input for the error detection model (Figure 3).

4.2 Error Detection Model

We used the output of NeckNet-18 (55 Blendshape and head motion parameters) while participants were watching stimulus videos and trained models that *detect when humans are reacting to failure*. We benchmarked this performance against other conventional data inputs, and tested the models’ ability to generalize to new participants

or within the same participant data. Figure 3 illustrates how study stages and datasets integrated into our processing methodology.

Problem Formulation: Detecting failure based on human reactions was formulated as a sequential decision-making problem: at each time step t , the environment is captured as a state variable $s_t \in S$ and the model output is an outcome label $f_t \in F$, where F is a discrete variable that describes the stimulus: 0 if *neutral*, 1 if *failure*.

Action Space: The reactions captured with NeckFace IR cameras and RGB video were labeled according to the type of stimulus video shown to participants. In videos where error occurs, only the moments following the exact failure moment were used as input data (and labeled 1, *reaction to failure*). Control videos, where there is no error, were used integrally, and labeled 0, *neutral state*. This reduced data imbalance across classes.

Datasets: We tested 4 data input sources, to test the performance of NeckFace against other commonly used methods.

- **NeckData:** this dataset was collected at 12 fps, consisting of 55 Blendshape and head motion parameters as described in Section 4.1, reconstructed using NeckNet-18.
- **OpenData:** Reactions to the stimulus videos were also recorded using a RGB camera at 30 fps. We ran OpenFace [3] on this data and obtained 49 Action Unit (AU), gaze and pose features. The dataset was undersampled to 12 fps, to match the NeckData timestamps.
- **NeckFaceIR:** We used video frames from NeckFace’s IR cameras to test direct output classification. Cameras collected data at 12 fps and this data was resized to 224x224. Post-processing and data augmentation were performed on each frame, according to ImageNet [21] statistics for data normalization.
- **RGBData:** We used frames from the RGB camera reaction videos, which were processed as in *NeckFaceIR* and downsampled to match the 12 fps sampling and timestamps.

We also tested normalizing the dataset and running a Principal Component (PC) Analysis for feature reduction. We kept only the PCs which explained 95% of the variance on the dataset, resulting in 15 (resp., 33) features for the NeckData (resp. OpenData) dataset. The total number of individual datapoints per dataset can be seen in Table 1. Because all participants watched the same stimulus videos, data is balanced across participants.

Models tested: To develop an error detection model from human reactions, we tested a diverse range of model architectures based on prior work in the space of robot error detection [30, 36, 42]. These included Recurrent Neural Networks (RNNs) such as Long Short-Term Memory (LSTM) networks [19], which excel at capturing long-term dependencies; Gated Recurrent Units (GRUs), known for their efficiency in training; and Bidirectional LSTMs (BiLSTMs), which process sequences in both directions. We also explored Transformer models, leveraging their self-attention mechanism for capturing global dependencies. Additionally, we implemented MiniRocket [10, 41, 42], a non-neural approach known for its speed and accuracy in time series classification. The gated Multi-Layer Perceptron (gMLP) [25] was tested as an alternative to attention-based models, using spatial gating units. Lastly, we included InceptionTime




Figure 4: NeckNet-18 architecture.

and InceptionTimePlus [17], an ensemble of deep Convolutional Neural Network (CNN) models designed specifically for time series classification. We also tested a similar Deep Neural Network to that suggested by Stiber et al. [39]—3 hidden layers, with 64, 128 and 64 units (multi-layer DNN, ml-DNN). For the video datasets, we used ResNet34 models [14], which are 34-layer CNNs pre-trained on ImageNet [21], and tested different model structures. ResNet34 is the base-model for the original version of NeckNet [8].

Evaluation Metrics: The models were evaluated based on the macro averages of the following metrics: accuracy, f1-score, precision and recall. We also considered margin-of-error metrics [9, 29, 36]—for a sample margin of size k , and for a sample i , the model prediction is considered correct if $y_{pred}^i \in [y_{pred}^{i-k}, y_{pred}^{i+k}]$. Similar to prior works, we consider these metrics to contemplate real-life error detection scenarios, where a small tolerance in timely detection does not invalidate the system.

Model Training: We implemented model training through tsai [28] and pytorch. For all model types and datasets, we performed hyperparameter tuning on a 70-20-10 train-val-test split, with 5 cross-validation folds with no overlapping participants, for 500 epochs. For time-series models, we implemented an interval length (i.e., number of samples in a training sample window) and stride (i.e., sliding length between sample windows) according to the methodology in Wachowiak et al. [41]. Labels were selected according to the mode (most frequent label) of each sample window. We picked the top-performing models based on average test accuracy across the 5 folds, and selected the best-performing epoch based on the same metric. For the Resnet34-based models (used on NeckFaceIR and RGBData datasets), best performing models were selected based on test accuracy on a single fold.

We additionally tested cross-training of the top-performing model types. That is, we used the best performing models on the *NeckData* and *OpenData* datasets, respectively, and fine-tuned the model’s last layer to the alternative dataset, following the same process for model selection as described above.

Finally, we tested single-participant generalization, based on minimum data required for good model performance. For this, we used n folds (n = number of participants), and trained the models on each i fold using a range between 5-45% of randomly selected data




Figure 5: Best performing model for NeckData, GRU_FCN.

Table 1: Samples (frames) per dataset per type of reaction.

Dataset	Neutral (0)	Error (1)	Total
NeckData	51745	37815	89560
OpenData	51580	37815	89395
NeckFaceIR	51745	37815	89560
RGBData	51578	37815	89393

for participant i . 20% of the data was used for validation and the remaining data was used for testing the model.

Final Dataset for Error Detection: The size of the datasets used for training are presented in Table 1. NeckData is obtained from reconstructing the Neckface IR camera data from the Stimulus round using NeckNet-18. The RGB-based datasets (RGBData and OpenData) are obtained through undersampling and sample matching to NeckData.

5 RESULTS

We present the main results from our paper. Our contributions are 1) developing a 3D facial reconstruction model, NeckNet-18 and 2) benchmarking error detection models trained on NeckFace-derived data against commonly used features (e.g. OpenFace).

5.1 3D Facial Mapping Model

Following the procedures described in Section 4.1, we ran 5-fold cross validation (non-overlapping temporal data) to test NeckNet-18 performance for 3D facial mapping. For Blendshape parameter data (facial motion), performance was $MAE_f = 34.0 \pm 6.9$ ($M \pm SD$), and $MAE_o = 7.0 \pm 1.3$ for head rotation angles. For reference, the original NeckNet performance [8] was $MAE_f = 25.6 \pm 5.1$, $MAE_o = 3.1$.

5.2 Error Detection Model

Table 2 presents the results from the best-performing models for each training method and dataset type. The best performing model on the NeckData dataset can be seen in Figure 5. *Pre-trained OpenData* (resp. NeckData) refers to the model pre-trained on NeckData (resp. OpenData) and fine-tuned in OpenData (resp. NeckData). All time-series models (NeckData and OpenData datasets) performed best in the non-normalized dataset, except for pre-trained and single

participant OpenData (PCA dataset). A full list of hyperparameters tested can be found in the study repository. In sum, our results for error detection from human facial cues are:

- *Time-series models* (trained on OpenFace features or Blendshapes): the model trained on NeckData outperforms the model trained using OpenFace features, by 5% accuracy and 10% F1.
- *Pre-trained Models* (trained on one dataset and fine-tuned on the other): performance is lower than single-dataset training. Model fine-tuned on NeckData outperforms model fine-tuned on OpenData, by 5% accuracy and 5% F1.
- *Single-Participant Models* (trained on subset of data from participant and tested on remaining data from that participant): for smaller training set (5% of data used from training, 95% for testing), NeckData model outperforms model trained using OpenFace features, by 6% accuracy and F1. For larger training set (45%), OpenData model slightly outperforms NeckData model, by 1% accuracy and F1. NeckData model achieves accuracy above 84% when trained on only 5% of data for a single participant.
- *Frame-Based Models* (trained on images, RGB videos from front-facing camera or videos from the IR cameras on the NeckFace device): model trained on NeckFace camera data outperforms RGB-trained ResNet34 by 1.5% accuracy but has lower F1 (2%).

6 DISCUSSION

Motivated by human social competence, where we continuously interpret external social cues from our peers to navigate social environments, our study explored a novel approach to more contextually-aware robots through error detection by leveraging human facial reactions captured through a neck-mounted device, NeckFace. Our contributions are twofold: 1) we adapted NeckNet [8] to a new dataset of human reactions, using a lighter model; 2) we developed models using data from the NeckFace device and NeckNet to detect errors from human reactions, finding good model performance both across and within participant data and outperforming conventional methodologies such as OpenFace.

6.1 3D Facial Mapping Model

The performance for NeckNet-18 was lower than that of the original NeckFace work [8]. This difference may be attributed to the following factors: 1) the original NeckFace work used more calibration data (7-minute rounds) while our calibration session had a duration of 5 minutes, of which only 4 minutes were initially used for training, 2) the original NeckFace device sampled at 30 fps while our device sampled at 12 fps. This is due to our use of a single Raspberry Pi device that records data from two cameras instead of using two Pi's. While implying a lower 3D mapping performance, these two adjustments were made to minimize the burden on participants regarding the duration of the session and device size, while preserving valid performance. According to Chen et al. [8], when MAE is under 40, the reconstructions should appear visually similar to the ground truth.

Table 2: Model hyperparameters and test set performance. Metrics are shown in $M \pm SD$ for 5 non-overlapping participant folds, or for all participants for single-participant models ($n = 25$). using data collected through the RGB camera are shown in grey rows. Percentage of participant data used for training is 5% for single-participant models.

Time-series Models					
Model	Dataset	Accuracy	Precision	Recall	F1
gMLP	OpenData	0.606 ± 0.051	0.563 ± 0.025	0.632 ± 0.072	0.535 ± 0.015
GRU_FCN	NeckData	0.658 ± 0.061	0.649 ± 0.060	0.667 ± 0.052	0.637 ± 0.070
Pre-trained Models					
gMLP	OpenData	0.569 ± 0.040	0.547 ± 0.028	0.572 ± 0.075	0.538 ± 0.019
gMLP	NeckData	0.619 ± 0.057	0.597 ± 0.042	0.619 ± 0.057	0.589 ± 0.047
Single-Participant Models					
InceptionTime	OpenData	0.788 ± 0.049	0.782 ± 0.048	0.785 ± 0.050	0.782 ± 0.049
GRU_FCN	NeckData	0.847 ± 0.048	0.840 ± 0.048	0.846 ± 0.050	0.842 ± 0.049
Frame-Based Models					
ResNet34	RGB	0.515 ± 0.025	0.505 ± 0.020	0.505 ± 0.018	0.493 ± 0.023
ResNet34	NeckFace_IR	0.530 ± 0.031	0.504 ± 0.027	0.501 ± 0.016	0.474 ± 0.014

6.2 Error Detection Model

We further expanded on prior work on robot error detection through human social cues by extensively exploring the potential of NeckFace-generated data in reaction classification. When comparing cross-participant performance, a GRU model trained on *NeckData* outperforms a gMLP model trained on *OpenData*, i.e. facial action units data extracted from video frames of the participants’ reactions recorded on a participant-facing camera, achieving a test accuracy of 65.8%. While low generalizability to unseen participants has been discussed [4], highlighting the need for larger datasets of human reactions to different error contexts, this performance is also higher than reported on previous works [30, 36]. Notably, an experimental factor which could impact this result is different levels of participant engagement. For example, Parreira et al. [30] collected data online through Prolific², wherein participants might be distracted or less engaged than when watching videos in-lab. Nonetheless, these results validate NeckFace as a device that effectively captures human reactions and NeckNet as a data reconstruction methodology in the context of this user study.

Interestingly, the pre-trained models (those trained on one dataset—*OpenData*, *NeckData*—and fine-tuned on the other) did not outperform single-dataset models. Still, we find that a model trained on OpenFace data and fine-tuned using NeckFace data outperforms a model trained on OpenFace data only. This indicates that the NeckFace approach may offer benefits even when utilizing transfer learning techniques. For frame-based models, interestingly, performance was comparable across the two datasets (video frames from an RGB participant-facing camera and from NeckFace’s IR camera), with both datasets generating models which perform slightly above chance.

The most striking results come from the single-participant models. Using just 5% of the data for training, the GRU model trained on NeckFace data achieved an accuracy of 84.7%, surpassing the InceptionTime model trained on OpenData (78.8%). This gap narrows when using 45% of the data for training, suggesting that NeckFace is most useful in contexts of few calibration data, allowing for effective personalized error detection systems. Further, the GRU model

used only has around 300,000 trainable parameters, which yields potential for lightweight computation and system portability.

The performance of NeckFace-based models, especially in single-participant scenarios, highlights the potential of this approach for creating more adaptable and personalized robot interaction systems. By capturing facial cues from a unique angle, NeckFace appears to provide rich, informative data that can be leveraged for error detection, namely through mapping reactions into a 3D plane rather than OpenFace’s 2-dimensional processing. By tapping into previously underutilized data sources, we can enhance robots’ ability to interpret human reactions and adjust their behavior accordingly.

6.3 General System Considerations

Detecting and responding to social errors is a critical component of creating socially competent robots that can seamlessly integrate into diverse human environments. Our use of human facial reactions as a signal for error detection builds directly from HRI literature [4, 31, 39]. While the technical approach leverages machine learning, the core objectives and implications of our work are deeply rooted in human-robot interaction applications. The inclusion of human and robot error videos in the stimulus dataset explores the potential of a generalized error detection system which would allow agents to contextualize failures in their environment. Future work may explore higher error granularity, e.g. through detecting robot error type [40]. To the best of our knowledge, no prior work has explored human reactions for multi-class error detection in the environment, which could improve robot’s social and functional performance.

We acknowledge that the proposed system implies additional steps and computation, including the calibration round and the 3D reconstruction using NeckNet-18 before error detection. This work aimed to push the boundaries of input data modalities used in HRI systems by exploring the NeckFace device and its functionalities. We verify that error detection accuracy in models trained on NeckData is higher by 12-15% than when directly using frame-based datasets. Similar results on frame-based models for non-overlapping participant testing had been reported [4, 30]. While these are non-exhaustive results and there is potential for RGB-based models to detect error effectively, higher computational demands for feature

²<https://www.prolific.com/>

extraction might be justified by better performing systems. This tradeoff warrants discussion, especially in more critical application scenarios.

6.4 Limitations and Future Work

It is important to highlight the scope and limitations of our work. The controlled environment of our user study may not fully represent the complexities of real-world human-robot interactions. For example, NeckFace is highly sensitive to lighting variations, which can compromise the performance of NeckNet. Further data collection in different environmental conditions, as well as additional research (e.g., in-person errors) is needed to validate the effectiveness of this approach across diverse populations and in various real-world scenarios, where noisy reactions are common.

The exclusion of participants from our dataset due to the inability to effectively transform the NeckFace camera data onto the 3D reconstruction brings about the tradeoff between larger datasets versus better-curated datasets. A longer calibration round or higher data frame rate could also provide more accurate NeckFace data reconstruction.

Finally, we note the ethical and sustainability implications of our work. While the envisioned use-case – more harmonious human-robot interactions, where robots play a collaborative and complementary role to human activities and needs, should contribute to increased human satisfaction, fulfillment, and efficiency, there is a risk of misuse of this technology. Broader discussions ought to be held to ensure frameworks where new tools are not misappropriated and contribute to a more just and equitable society. Related to this, access to technology is also an important indicator of development and equity. The NeckFace device is relatively inexpensive and easy to replicate, and while it requires high power consumption in its current version, newer and more efficient materials should mitigate these limitations.

7 CONCLUSION

This research underscores our claim that expanding human-in-the-loop robot sensing can lead to more adaptable and personalized robot interaction systems. Leveraging the social competence humans exhibit through interpreting social cues, NeckFace offers a new perspective on capturing these previously untapped data sources. The performance of NeckFace-derived models, especially in single-participant scenarios where personalization is key, reinforces its potential for creating more context-aware robots. By effectively capturing facial cues, NeckFace could enable robots to better understand and respond to human reactions in real-time. This capability is essential for the integration into dynamic human environments, pushing robots beyond simple task execution toward becoming socially aware collaborators guided by human behavior. The NeckFace form factor and classification results give confidence for further exploration into other HRI domains, where traditional methods have previously been limited.

REFERENCES

- [1] R. M. Aronson. Gaze for error detection during human-robot shared manipulation. In *RSS Workshop: Towards a Framework for Joint Action*, 2018.
- [2] T. Baltrušaitis, P. Robinson, and L.-P. Morency. Openface: an open source facial behavior analysis toolkit. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–10. IEEE, 2016.
- [3] T. Baltrušaitis, A. Zadeh, Y. C. Lim, and L.-P. Morency. Openface 2.0: Facial behavior analysis toolkit. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 59–66, 2018. doi: 10.1109/FG.2018.00019.
- [4] A. Bremers, M. T. Parreira, X. Fang, N. Friedman, A. Ramirez-Aristizabal, A. Pabst, M. Spasojevic, M. Kuniavsky, and W. Ju. The bystander affect detection (bad) dataset for failure detection in hri, 2023.
- [5] A. Bremers, A. Pabst, M. T. Parreira, and W. Ju. Using social cues to recognize task failures for hri: Overview, state-of-the-art, and future directions, 2024. URL <https://arxiv.org/abs/2301.11972>.
- [6] D. E. Cahya, R. Ramakrishnan, and M. Giuliani. Static and temporal differences in social signals between error-free and erroneous situations in human-robot collaboration. In *Social Robotics: 11th International Conference, ICSR 2019, Madrid, Spain, November 26–29, 2019, Proceedings*, page 189–199, Berlin, Heidelberg, 2019. Springer-Verlag. ISBN 978-3-030-35887-7. doi: 10.1007/978-3-030-35888-4_18. URL https://doi.org/10.1007/978-3-030-35888-4_18.
- [7] K. Candon, N. C. Georgiou, H. Zhou, S. Richardson, Q. Zhang, B. Scassellati, and M. Vázquez. React: Two datasets for analyzing both human reactions and evaluative feedback to robots over time. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction, HRI '24*, page 885–889, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9789400703225. doi: 10.1145/3610977.3637480. URL <https://doi.org/10.1145/3610977.3637480>.
- [8] T. Chen, Y. Li, S. Tao, H. Lim, M. Sakashita, R. Zhang, F. Guimbretière, and C. Zhang. Neckface: Continuously tracking full facial expressions on neck-mounted wearables. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 5(2), Jun 2021. doi: 10.1145/3463511. URL <https://doi.org/10.1145/3463511>.
- [9] I. de Kok and D. K. Heylen. A survey on evaluation metrics for backchannel prediction models. In *Interdisciplinary Workshop on Feedback Behaviors in Dialog, Stevenson, Washington, USA: Proceedings of the Interdisciplinary Workshop on Feedback Behaviors in Dialog*, pages 15–18. University of Texas, 2012.
- [10] A. Dempster, D. F. Schmidt, and G. I. Webb. Minirocket: A very fast (almost) deterministic transform for time series classification. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. ACM, 2021. doi: 10.1145/3447548.3467231. URL <http://dx.doi.org/10.1145/3447548.3467231>.
- [11] M. A. Diaz, M. Voß, A. Dillen, B. Tassignon, L. Flynn, J. Geeroms, R. Meeusen, T. Verstraten, J. Babić, P. Beckerle, and K. De Pauw. Human-in-the-loop optimization of wearable robotic devices to improve human-robot interaction: A systematic review. *IEEE Transactions on Cybernetics*, 53(12):7483–7496, 2023. doi: 10.1109/TCYB.2022.3224895.
- [12] M. Giuliani, N. Mirnig, G. Stollnberger, S. Stadler, R. Buchner, and M. Tscheligi. Systematic analysis of video data from different human-robot interaction studies: a categorization of social signals during error situations. *Frontiers in psychology*, 6:931, 2015.
- [13] A. Gruebler and K. Suzuki. Measurement of distal emg signals using a wearable device for reading facial expressions. In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, pages 4594–4597. IEEE, 2010.
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition, 2015. URL <https://arxiv.org/abs/1512.03385>.
- [15] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [16] K.-S. Hwang, J. Ling, Y.-Y. Chen, and W.-H. Wang. Reward shaping for reinforcement learning by emotion expressions. In *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 1288–1293, New York, NY, USA, 2014. IEEE, IEEE.
- [17] H. Ismail Fawaz, B. Lucas, G. Forestier, C. Pelletier, D. F. Schmidt, J. Weber, G. I. Webb, L. Idoumghar, P.-A. Muller, and F. Petitjean. Inceptiontime: Finding alexnet for time series classification. *Data Mining and Knowledge Discovery*, 34(6):1936–1962, Sept. 2020. ISSN 1573-756X. doi: 10.1007/s10618-020-00710-y. URL <http://dx.doi.org/10.1007/s10618-020-00710-y>.
- [18] C. Junge, P. M. Valkenburg, M. Deković, and S. Branje. The building blocks of social competence: Contributions of the consortium of individual development. *Developmental Cognitive Neuroscience*, 45:100861, 2020. ISSN 1878-9293. doi: <https://doi.org/10.1016/j.dcn.2020.100861>. URL <https://www.sciencedirect.com/science/article/pii/S1878929320301110>.
- [19] F. Karim, S. Majumdar, H. Darabi, and S. Chen. Lstm fully convolutional networks for time series classification. *IEEE Access*, 6:1662–1669, 2018. doi: 10.1109/ACCESS.2017.2779939.
- [20] D. Kontogiorgos, M. Tran, J. Gustafson, and M. Soleymani. A systematic cross-corpus analysis of human reactions to robot conversational failures. In *Proceedings of the 2021 International Conference on Multimodal Interaction*, pages 112–120, 2021.
- [21] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.
- [22] K. Li, R. Zhang, B. Liang, F. Guimbretière, and C. Zhang. Eario: A low-power acoustic sensing earable for continuously tracking detailed facial movements. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 6(2), jul 2022. doi: 10.1145/3534621. URL <https://doi.org/10.1145/3534621>.

[23] K. Li, R. Zhang, S. Chen, B. Chen, M. Sakashita, F. Guimbretiere, and C. Zhang. Eyeecho: Continuous and low-power facial expression tracking on glasses. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, CHI '24, New York, NY, USA, 2024*. Association for Computing Machinery. ISBN 9798400703300. doi: 10.1145/3613904.3642613. URL <https://doi.org/10.1145/3613904.3642613>.

[24] J. Lin, Z. Ma, R. Gomez, K. Nakamura, B. He, and G. Li. A review on interactive reinforcement learning from human social feedback. *IEEE Access*, 8:120757–120765, 2020. doi: 10.1109/ACCESS.2020.3006254.

[25] H. Liu, Z. Dai, D. R. So, and Q. V. Le. Pay attention to mlps, 2021. URL <https://arxiv.org/abs/2105.08050>.

[26] F. Loureiro, J. a. Avelino, P. Moreno, and A. Bernardino. Self-perception of interaction errors through human non-verbal feedback and robot context. In *Social Robotics: 14th International Conference, ICSR 2022, Florence, Italy, December 13–16, 2022, Proceedings, Part II*, page 475–487, Berlin, Heidelberg, 2023. Springer-Verlag. ISBN 978-3-031-24669-2. doi: 10.1007/978-3-031-24670-8_42. URL https://doi.org/10.1007/978-3-031-24670-8_42.

[27] C. Lugaressi, J. Tang, H. Nash, C. McClanahan, E. Uboweja, M. Hays, F. Zhang, C.-L. Chang, M. G. Yong, J. Lee, et al. Mediapipe: A framework for building perception pipelines. *arXiv preprint arXiv:1906.08172*, 2019.

[28] I. Oguiza. tsai - a state-of-the-art deep learning library for time series and sequential data. Github, 2023. URL <https://github.com/timeseriesAI/tsai>.

[29] M. T. Parreira, S. Gillet, and I. Leite. Robot duck debugging: Can attentive listening improve problem solving? In *Proceedings of the 25th International Conference on Multimodal Interaction*, pages 527–536, 2023.

[30] M. T. Parreira, S. G. Lingaraju, A. Ramirez-Aristizabal, M. Saha, M. Kuniavsky, and W. Ju. A study on domain generalization for failure detection through human reactions in hri. *arXiv preprint arXiv:2403.06315*, 2024.

[31] M. T. Parreira, S. G. Lingaraju, A. Ramirez-Aristizabal, A. Bremers, M. Saha, M. Kuniavsky, and W. Ju. “bad idea, right?” exploring anticipatory human reactions for outcome prediction in hri. In *2024 33rd IEEE International Conference on Robot and Human Interactive Communication (ROMAN)*, pages 2072–2078, 2024. doi: 10.1109/RO-MAN60168.2024.10731310.

[32] V. Rantanen, H. Venesvirta, O. Spakov, J. Verho, A. Vetek, V. Surakka, and J. Lekkala. Capacitive measurement of facial activity intensity. *IEEE Sensors journal*, 13(11):4329–4338, 2013.

[33] J. Ravishankar, M. Doering, and T. Kanda. Zero-shot learning to enable error awareness in data-driven hri. In *2024 19th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 592–601, 2024.

[34] C. Reverdy, S. Gibet, and C. Larboulette. Optimal marker set for motion capture of dynamical facial expressions. In *Proceedings of the 8th ACM SIGGRAPH Conference on Motion in Games, MIG '15*, page 31–36, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450383424. doi: 10.1145/3447993.3483252.

for Computing Machinery. ISBN 9781450339919. doi: 10.1145/2822013.2822042. URL <https://doi.org/10.1145/2822013.2822042>.

[35] P. Slade, C. Atkeson, J. M. Donelan, H. Houdijk, K. A. Ingraham, M. Kim, K. Kong, K. L. Poggensee, R. Riener, M. Steinert, J. Zhang, S. H. Collins, H. John, and A. Paulson. On human-in-the-loop optimization of human–robot interaction. *Nature* 2024 633:8031, 633:779–788, Sep 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07697-2. URL <https://www.nature.com/articles/s41586-024-07697-2>.

[36] M. Spitale, M. T. Parreira, M. Stiber, M. Axelsson, N. Kara, G. Kankariya, C.-M. Huang, M. Jung, W. Ju, and H. Gunes. Err@hri 2024 challenge: Multimodal detection of errors and failures in human–robot interactions, 2024. URL <https://arxiv.org/abs/2407.06094>.

[37] M. Stiber. Effective human–robot collaboration via generalized robot error management using natural human responses. In *Proceedings of the 2022 International Conference on Multimodal Interaction, ICMI '22*, page 673–678, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450393904. doi: 10.1145/3536221.3557028. URL <https://doi.org/10.1145/3536221.3557028>.

[38] M. Stiber and C.-M. Huang. Not all errors are created equal: Exploring human responses to robot errors with varying severity. In *Companion Publication of the 2020 International Conference on Multimodal Interaction, ICMI '20 Companion*, page 97–101, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450380027. doi: 10.1145/3395035.3425245. URL <https://doi.org/10.1145/3395035.3425245>.

[39] M. Stiber, R. H. Taylor, and C.-M. Huang. On using social signals to enable flexible error-aware hri. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction, HRI '23*, page 222–230, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450399647. doi: 10.1145/3568162.3576990. URL <https://doi.org/10.1145/3568162.3576990>.

[40] P. Trung, M. Giuliani, M. Miksch, G. Stollnberger, S. Stadler, N. Mirlig, and M. Tscheligi. Head and shoulders: Automatic error detection in human–robot interaction. New York, NY, USA, 2017. ACM. ISBN 9781450355438.

[41] L. Wachowiak, P. Tisnikar, G. Canal, A. Coles, M. Leonetti, and O. Celiktutan. Predicting when and what to explain from multimodal eye tracking and task signals. *IEEE Transactions on Affective Computing*, (01):1–12, jun 2024. ISSN 1949-3045. doi: 10.1109/TAFFC.2024.3419696.

[42] L. Wachowiak, P. Tisnikar, A. Coles, G. Canal, and O. Celiktutan. *A Time Series Classification Pipeline for Detecting Interaction Ruptures in HRI Based on User Reactions*. ACM, Aug. 2024. doi: 10.1145/3678957.3688386.

[43] Y. Wu, V. Kakaraparthi, Z. Li, T. Pham, J. Liu, and P. Nguyen. Bioface-3d: continuous 3d facial reconstruction through lightweight single-ear biosensors. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking, MobiCom '21*, page 350–363, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383424. doi: 10.1145/3447993.3483252. URL <https://doi.org/10.1145/3447993.3483252>.