



# ERR@HRI 2.0 Challenge: Multimodal Detection of Errors and Failures in Human-Robot Conversations

Shiye Cao  
Johns Hopkins University  
Baltimore, USA  
shiyecao@cs.jhu.edu

Maia Stiber  
Microsoft Research  
Redmond, USA  
maia.stiber@microsoft.com

Amama Mahmood  
Johns Hopkins University  
Baltimore, USA  
amama.mahmood@jhu.edu

Maria Teresa Parreira  
Cornell University  
Ithaca, USA  
mb2554@cornell.edu

Wendy Ju  
Cornell Tech  
New York, USA  
wendyju@cornell.edu

Micol Spitale  
Politecnico di Milano  
Milano, Italy  
micol.spitale@polimi.it

Hatice Gunes  
University of Cambridge  
Cambridge, UK  
hatice.gunes@cl.cam.ac.uk

Chien-Ming Huang  
Johns Hopkins University  
Baltimore, USA  
chienming.huang@jhu.edu

## Abstract

The integration of large language models (LLMs) into conversational robots has made human-robot conversations more dynamic. Yet, LLM-powered conversational robots remain prone to errors, e.g., misunderstanding user intent, prematurely interrupting users, or failing to respond altogether. Detecting and addressing these failures is critical for preventing conversational breakdowns, avoiding task disruptions, and sustaining user trust. To tackle this problem, the ERR@HRI 2.0 Challenge provides a multimodal dataset of LLM-powered conversational robot failures during human-robot conversations and encourages researchers to benchmark machine learning models designed to detect robot failures. The dataset includes 16 hours of dyadic human-robot interactions, incorporating facial, speech, and head movement features. Each interaction is annotated with the presence or absence of robot errors from the system perspective, and perceived user intention to correct for a mismatch between robot behavior and user expectation. Participants are invited to form teams and develop machine learning models that detect these failures using multimodal data. Submissions will be evaluated using various performance metrics, including detection accuracy and false positive rate. This challenge represents another key step toward improving failure detection in human-robot interaction through social signal analysis.

## CCS Concepts

• Computer systems organization → Robotics; • Computing methodologies → Artificial intelligence; • Human-centered computing → Human computer interaction (HCI).

## Keywords

human-robot interaction, failure detection, multimodal behavior, conversational agents

### ACM Reference Format:

Shiye Cao, Maia Stiber, Amama Mahmood, Maria Teresa Parreira, Wendy Ju, Micol Spitale, Hatice Gunes, and Chien-Ming Huang. 2025. ERR@HRI 2.0 Challenge: Multimodal Detection of Errors and Failures in Human-Robot Conversations. In *Proceedings of the 33rd ACM International Conference on Multimedia (MM '25)*, October 27–31, 2025, Dublin, Ireland. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3746027.3762073>

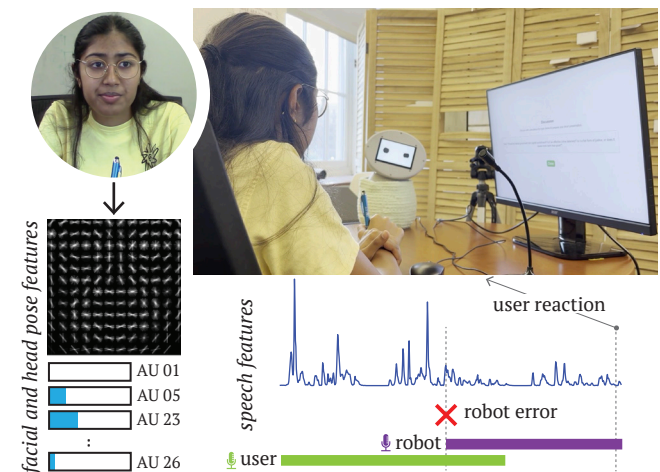


Figure 1: Illustration of the interaction and data recording setup, together with extracted features.

## 1 Introduction

In recent years, the integration of large language models (LLMs) into conversational robots—embodied media—has enabled more natural human-robot conversations. While LLMs can mitigate certain



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

MM '25, Dublin, Ireland

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2035-2/2025/10

<https://doi.org/10.1145/3746027.3762073>

speech recognition errors that often cause traditional conversational robots to fail, these advanced systems remain vulnerable to issues such as interrupting users and misinterpreting user intent [8, 20, 21]. Such failures can lead to conversational breakdowns, disrupt interaction flow, and negatively affect users' perception of the robot. As LLM-powered conversational robots are deployed in real-world scenarios, the need to understand and mitigate their failures is a pressing issue. Prior work has discovered some strategies to help robots rebuild user trust and recover from errors [1, 10, 17]; however, the application of these strategies hinges on the robot's ability to detect its own errors effectively [34]. Status quo robot error detection methods often either depend on task-specific or domain-specific information—which can be hard to generalize across people, tasks, and errors—or rely on users to detect and report errors—which can be delayed. Recent works showed the possibility of detecting robot error through leveraging social signals exhibited by user reactions and responses to robot errors in multimedia streams [4, 16, 24, 29, 33, 38]. To further explore this potential and benchmark multimodal robot failure detection in HRI, the ERR@HRI 2.0 challenge focuses on **multimodal detection of robot errors and failures in dyadic human-robot conversations**.

The ERR@HRI initiative serves as a platform for benchmarking not only HRI datasets but also machine learning models aimed at detecting robot errors and failures from multimedia. The 2024 ERR@HRI challenge, held as part of the 2024 ACM International Conference on Multimodal Interaction (ICMI'24), focused on modeling robot errors using a multimodal dataset collected in a real-world setting where a robotic coach delivered well-being coaching practices to each individual separately over four weeks [30]. Building on this, the ERR@HRI 2.0 challenge uses a new multimodal dataset (see Fig. 1) comprised of multimodal behavioral signals extracted from audio and video recordings of human conversations with two distinctive LLM-powered embodied media: a social robot with more anthropomorphic features and expressions, and a smart speaker. The user and the robot collaborated on five tasks during the conversations, ranging from medical self-diagnosis, trip planning, selecting items more beneficial during a desert survival simulation, to discussing whether the federal government should ban capital punishment and whether universities should have their own police force. Additionally, the ERR@HRI 2.0 challenge is also unique in that it contains two sub-challenges to tackle robot error detection from two perspectives. Sub-challenge 1 targets the detection of robot errors and failures from the system perspective, while sub-challenge 2 targets the detection of robot errors and failures from the user perspective.

Through this challenge, we hope to bring together researchers and practitioners from the multimedia, robotics, and human-robot interaction communities, while highlighting the underlying multimedia aspects of the problem. The techniques developed through this challenge will contribute to advancing our understanding of how multimedia can enhance the analysis and interpretation of interactions between humans and autonomous robots.

## 2 Related Work

### 2.1 Robot Errors

As robots become more capable, they are increasingly being integrated into daily life, assisting and collaborating with human users across a range of tasks and contexts [12, 19, 23, 27]. However, errors and failures remain common and often inevitable in robots deployed in real-world settings [9, 14, 22]. For a robot to be truly effective in a human-robot team, it must not only complete its assigned tasks but also adapt to the user and align its behavior with the user's expectations—thereby enhancing team performance and user satisfaction [35]. This highlights a limitation in existing robot evaluation approaches that only assess the robot based on task completion. Instead, we need to consider how the task was performed and whether the robot's behavior conformed to their human teammate's expectations. Therefore, in this work, we adopt a broader definition of robot errors to assess the robot behaviors from two distinct perspectives: 1) *robot errors from the system perspective*—deviations in the robot behavior from its expected behavior by system design; and 2) *robot errors from the user perspective*—deviations in the robot behavior from the user's mental model of the expected robot behavior.

### 2.2 Robot Error Detection Using Social Signals

Social signals are multimodal behavioral cues (both verbal and non-verbal) that can convey information about one's emotions, intentions, attitudes, and social dynamics within interactions [6, 7, 26, 31, 36]. In human-robot interaction (HRI), users' natural responses (social signals) during the interaction can impart information about their own internal state (e.g., user uncertainty) and their mental model of the robot [15, 28, 32, 35, 37]. As such, social signals exhibited by users have been used to aid collaboration in HRI by conveying user preferences (e.g., [6]), user's need for help (e.g., [39]), and engagement breakdowns (e.g., [3]).

Social signals have also been used to facilitate automatic error detection in HRI. Robot errors naturally elicit responses from users [22]. Capitalizing on those elicitations expressed as multimodal streams, recent works have shown the possibility of using them for detecting robot errors [4, 13, 18, 29, 30, 32–34]. The 2024 ERR@HRI challenge [30] at the 2024 ACM International Conference on Multimodal Interaction focused on using a multimodal dataset collected in a real-world setting where a robotic coach delivered well-being coaching practices to each individual separately over four weeks. In an effort to further explore the potential of this and to benchmark multimodal robot failure detection in HRI, the ERR@HRI 2.0 challenge focuses on the multimodal detection of robot errors and failures in dyadic human-robot conversations across two unique embodiments and five tasks.

## 3 The ERR@HRI 2.0 Challenge

In continuation of previous year's ERR@HRI [30] efforts to encourage researchers to develop and benchmark multimodal machine learning models to detect robot errors during human-robot interaction, the ERR@HRI 2.0 challenge invites participants to develop multimodal ML models designed to detect conversational robot failures using a dataset that includes facial, head pose, and speech

**Table 1: Dataset characteristics**

Embodiments	Tasks	Time	# of Sessions	# of Robot Errors	
				System Perspective	User Perspective
Social Robot [8]	Survival	2 hr	22	73	29
	Discussion	1 hr 43 min	20	85	82
Voice Assistant[21]	Medical Self-Diagnosis	3 hr 50 min	20	72	40
	Trip Planning	4 hr 41 min	20	96	46
	Discussion	3 hr 36 min	19	88	27
	<b>Total</b>	<b>15 hr 50 min</b>	<b>101</b>	<b>414</b>	<b>224</b>

features extracted from 16 hours of dyadic human-robot conversational media.

### 3.1 Definition of Robot Errors

In this work, we are interested in detecting robot errors both from the system perspective and the user perspective (Section 2). These two types of robot errors overlap but are distinctive as deviations from the system’s designed behavior may not always constitute errors from the user’s perspective. Similarly, system’s designed behavior that deviates from user’s expected system behaviors are not always errors from the system’s perspective but are considered errors or failures from the user’s perspective. Therefore, it is important to consider both user perceived errors and system perceived errors when creating models that detect errors during human-robot interaction. We operationalize *robot errors from the system perspective* as a noticeable misalignment between robot behavior and its expected behavior by system design. We operationalized *robot errors from the user perspective* as noticeable verbal and non-verbal user-initiated disruptive interruptions that are intended to disrupt and correct for a misalignment between the robot behavior and the user’s expected robot behavior. We want to emphasize that our operationalized definition of robot error used in this work does not capture all robot errors from the system perspective nor the user perspective. However, this definition allows for robot errors to be labeled from a third-party perspective, without having the participants self-report when they felt like the robot made a mistake. A similar approach was used in prior work to define other user internal states, *e.g.*, confusion [31].

### 3.2 The Tasks

Hence, the ERR@HRI 2.0 Challenge consists of two sub-challenges:

- (1) Sub-challenge 1. Detection of robot errors from the system’s perspective
- (2) Sub-challenge 2. Detection of robot errors from the user’s perspective

In sub-challenge 1, we operationalize robot error from the system’s perspective as deviations in the robot behavior from its designed behavior. Examples of this type of robot errors include user intention recognition errors, interrupting the user, and not responding to the user.

In sub-challenge 2, we operationalize robot error from the user’s perspective as detection of user intention to correct for a mismatch between robot behavior and their expectation (*i.e.*, user-initiated disruptive interruptions). Examples of this include user-initiated

verbal and non-verbal interruptions that are perceived to be disruptive in intention.

### 3.3 Dataset

The dataset consists of data from 42 users (21 female, 21 male), in a total of 101 sessions and roughly 950 minutes (15 hours and 50 minutes 48 seconds) of interaction collected in a previous dyadic human-social robot conversation study [8] and dyadic human-voice assistant conversation study [21]. Table 1 summarizes additional characteristics of the dataset.

The human-social robot conversational data consists of 3 hours and 43 minutes 50 seconds of audio and video data of 22 unique participants interacting with a social robot to complete a set of one to two tasks including selected items to aid survival in a desert survival simulation (1 hour 43 minutes 11 seconds) and discussing whether the federal government should ban capital punishment (2 hours 39 seconds). For more details on the task and the study condition under which the data was collected, please refer to [8].

The human-voice assistant conversational data consists of 12 hours and 6 minutes 58 seconds of audio and video data, collected from 20 unique participants interacting with the voice assistant on a set of one to three tasks including medical self-diagnosis (3 hours 50 minutes 29 seconds), planning a day trip in Edinburgh (4 hours 40 minutes 45 seconds), and discussing whether universities should have their own police forces (3 hours 35 minutes 44 seconds). For more details on the task and the study condition under which the data was collected, please refer to the paper [21].

**3.3.1 Feature Extraction.** We took the audio (user speech and robot speech) and video (camera facing user faces) recordings of the interactions from the two studies and used off-the-shelf state-of-the-art methods to extract facial features, head pose features, audio features, and transcribed speech features:

- (1) Facial and head pose features: We used the OpenFace 2.2.0 toolkit [2] to identify facial expressions by detecting the presence of 18 facial action units (AUs) and estimating the intensity levels for 17 of those AUs. We also use the OpenFace 2.2.0 toolkit to estimate the location (x, y, and z coordinates) and rotation (x, y, and z coordinates) of the user’s head. These features combined with the toolkit output on its confidence and success results in a total of 43 features per frame at a rate of 30 frames per second.
- (2) Audio features: We used the openSMILE toolbox [11] to extract 25 audio features from frames of 20ms length every 10ms using the eGeMAPSv02 feature set. The eGeMAPSv02

feature set contains loudness, alpha ratio, Hammarberg index, spectral slope between 0 and 500 Hz, spectral slope between 500 and 1500 Hz, spectral flux, first four Mel Frequency Cepstral Coefficients (MFCCs 1–4), F0, jitter local, shimmer local, harmonics-to-noise ratio, log energy ratio of F0, frequency, bandwidth, and log amplitude relative to F0 of F1, F2, and F3.

- (3) **Transcribed speech features:** We used Google’s speaker diarization to detect user and robot speaker turns in audio recordings. Based on the result, we extracted 515 features showing when the robot/user is speaking, the word count of their speech content, the speech-to-text confidence, and created a 512-dimension embedding of the transcribed speech using a ViT-B/16 Transformer architecture Contrastive Language–Image Pre-training (CLIP) [25] embedding.

To preserve user privacy, we only released the 581 multimodal behavioral features extracted from the recordings for modeling.

**3.3.2 Labels.** Two coders annotated the videos using the Datavyu video annotation tool. The coders first familiarized themselves with system design and expected behavior of the robot by design. Then, they marked the start and end time of observable deviation of robot behavior from its designed behavior, start and end time of any observable user reactions to the robot error from the system behavior (if any), and the start and end time of any observable user verbal and non-verbal disruptive interruption intended to correct for a mismatch between robot behavior and their expectation. The two coders first independently annotated ten sessions (two sessions from each task). Then, they discussed to resolve conflicts in the annotations from the ten sessions. They reached 100% agreement after the discussion and proceeded to annotate the rest of the data.

The labels were defined as follows:

- (1) **Robot error from system perspective:** The robot makes a mistake such as interrupting or not responding to the user, or responding with an error message or an utterance that is not appropriate for what the user has just said.
- (2) **Reaction to robot error:** Observable verbal and non-verbal user reaction to the *robot error from the system perspective*. Users did not react to every error.
- (3) **Robot error from user perspective:** The user displays behavior (verbal or non-verbal) that signals an intention to correct for a mismatch between robot behavior and their expectation such as user-initiated disruptive interruption. Disruptive interruption is defined as when the listener challenges the speaker’s control and disrupts the conversational flow to take the floor, change the subject, or avoid unwanted information. This behavior suggests that there exists some mismatch between the user’s expectation for the robot and the robot’s behavior.

We added annotations for observable user reaction to robot error from the system perspective as an additional potential label to facilitate training models to detect robot errors from the system perspective for sub-challenge 1. The dataset contained a total of 414 errors from the system perspective and 224 errors from the user perspective. See Table 1 for the distribution of the labels across the different tasks.

**3.3.3 Training and Test Sets.** We divided the ERR@HRI 2.0 dataset into a training set and test sets by splitting the dataset using a subject-independent strategy (*i.e.*, the training, validation, and testing sets do not include data from the same subjects). We randomly selected 20 sessions from 5 users to form the test set. This resulted in a training set composed of 81 interactions (47 with voice assistants and 34 with social robots); a test set composed of 20 interactions (12 with voice assistants and 8 with social robots).

### 3.4 Metrics

We assess the model’s performance on each of the two sub-challenges in two ways: 1) the model’s error detection performance over fixed, pre-segmented windows, a simplified *offline* evaluation scheme compared to the streaming scenario; 2) the model’s error detection performance *on-the-fly*, simulating how the model would perform in a streaming scenario. We use different metrics for *offline* evaluation and *on-the-fly* evaluation due to the difference in evaluation goals.

We define an error as detected when the window in which the error is predicted overlaps with the ground truth error duration labeled. Following this definition, a prediction is a *true positive* when it overlaps with the ground truth error/reaction windows. We added a tolerance of one second to the start and end labels to account for label inaccuracies. If the two windows overlap by any amount then that is considered a true positive. Following this logic, we consider a prediction to be a *false positive*, when the prediction window does not overlap with any ground truth windows. We used these definitions to evaluate overall model performance.

In summary, we used the following metrics to evaluate the trained models in this work:

#### **Offline evaluation metrics:**

- Area Under the Receiver Operating Characteristic Curve (AUC)
- F1 score (F1)
- Accuracy (Acc)
- Balanced Accuracy Score (UAR) [5]

#### **On-the-fly evaluation metrics:**

- Percentage of errors/reactions detected (%Detected): This metric evaluates the detection performance of the models. It is calculated as the number of true positives (as defined above) detected divided by the number of errors/reactions in the ground truth.
- Number of false positives (#FPs): This metric measures the number of false positives.
- Overall F1 score (overall F1): This metric offers a balanced evaluation of detection and false positives.

### 3.5 Evaluation

Challenge participants were given access to the training data and features to develop their ML models. Then, they were asked to submit models and weights, and the organizers evaluated the submitted models on the test set (the test set was released to the challenge participants without labels one month prior to the submission deadline). To prevent overfitting on the test set, each participating group was allowed to submit their models and results for the test set up to three times.

As participating teams may use different training and data processing techniques, we only considered the overall F1-score when comparing and ranking performance across teams. This metric simulates the model performance in a streaming setting and requires the model have a balance between detecting robot errors (true positives) and false positive. This provided additional flexibility in the choice of training mechanisms in the teams. Metrics were calculated using the same script provided to participants in the study repository. Challenge participants were also asked to submit a paper describing their model, and their works were peer-reviewed by two reviewers.

### 3.6 Baseline Models

We provided a multimodal baseline for each of the sub-challenges<sup>1</sup>.

**3.6.1 Data Pre-Processing.** To facilitate training, we provided code that creates windowed chunks from the feature streams. A windowed chunk is considered to contain an error if it overlaps with the interval that is labeled robot error. We experimented with a range of window and step-sizes and decided on a window size of 3 seconds and step size of 0.5 seconds. Participating teams were encouraged to develop their own data pre-processing practices during models and were welcome to use different window sizes and step sizes. However, to prevent teams from using extremely large window sizes, which would result in significant delays in detection if these models were to be run in real-time, we limited the maximum allowed window size to 12 seconds for sub-challenge 1 and five seconds for sub-challenge 2. This maximum allowed window size is determined from the dataset and is set to be the average length of robot error/user reaction for that sub-challenge.

We chose to use a window of three seconds with a step size of 0.5 seconds. We inputted any missing entries in the training features with the corresponding feature’s mean. Since robot errors were sparse, this created a class imbalance, with only 11.35% of the windows with robot errors for sub-challenge 1 and 3.33% for sub-challenge 2.

**3.6.2 Training.** Due to the data imbalance between windows with robot errors and windows without robot errors, we applied synthetic minority over-sampling technique (SMOTE) to perform over-sampling to balance the training data. SMOTE synthesizes new minority-class examples in feature space to balance the classes, reducing bias toward the no-error class.

We explored a set of standard modeling approaches—random forest (RF), explainable boosting machine (EBM), Extreme Gradient Boosting (XGBoost), long short-term memory (LSTM), and single-layer transformers—while leaving room for participants to innovate their approaches for detection and classification. We decided to use random forest as it performed best out of these during our exploration.

We used leave-one-out cross validation and performed hyperparameter tuning (we tuned the number of estimators, max depth, min samples split, min samples leaf, max features, and class weight) using AUC on the validation set as the metric to pick the best hyperparameters. In the end, we trained a random forest using the best hyperparameters on all the training data and predicted on the

**Table 2: Baseline models’ window-based performance for the two sub-challenges.**

	AUC	F1	Acc	UAR
<b>Sub-challenge 1</b>	0.59	0.49	0.83	0.50
<b>Sub-challenge 2</b>	0.70	0.51	0.97	0.51

**Table 3: Baseline models’ on-the-fly performance for the two sub-challenges.**

	%Detected	#FPs	Overall F1
<b>Sub-challenge 1</b>	39.54%	163	0.24
<b>Sub-challenge 2</b>	29.03%	97	0.13

held-out test set. In line with the process for challenge participants, we adjusted the model based on feedback on the model performance on the test set. We adjusted the predicted threshold (sub-challenge 1: 0.55 and sub-challenge 2: 0.45) to reduce the number of false positives in sub-challenge 1 and increase the number of true positives in sub-challenge 2. The resulting model had an overall F1 score of 0.24 in sub-challenge 1 and an F1 score of 0.29 in sub-challenge 2. The baseline models’ performance on window-based evaluation metrics is shown in Table 2 and the models’ performance on the on-the-fly evaluation metrics is shown in Table 3

### 3.7 Participation and Conclusion

This paper introduces the ERR@HRI 2.0 Challenge organized in conjunction with the ACM International Conference on Multimedia 2025 (ACM-MM’25), which focuses on detecting robot errors and failures in human-robot conversations. A total of nine teams registered to participate in the challenge. Two teams submitted models to each of the sub-challenges and surpassed the baseline model performance. The two teams will be invited to submit a workshop-style paper describing their ML solutions and results on the dataset, as well as a publicly available code repository. We hope to continue hosting further ERR@HRI Grand Challenges in the coming years utilizing diverse human-robot interaction multimedia datasets. Future challenges should explore multimodal detection of different types of errors (physical, cognitive), severity of errors, types of robots (physical manipulator, mobile robots) in different settings (in-the-wild) to further advance multimodal detection of robot errors and failures.

### Acknowledgments

S. Cao, M. Stiber, A. Mahmood & C. M. Huang’s work was supported by the National Science Foundation award #2141335 and National Science Foundation award #2143704. M. Spitale has been supported by PNRR-PE-AI FAIR project funded by the NextGeneration EU program. H. Gunes has been supported by the EPSRC project ARoEQ under grant ref. EP/R030782/1. This paper has been proofread by a language model (AI), and the authors have read through the resulting content to ensure it accurately reflects the original intent. We would like to acknowledge Yifan Xu for their assistance with the baseline model training and Nithish Krishna Shreenevasan for their contribution to the data labeling.

<sup>1</sup>Baseline code and models: <https://github.com/ERR-HRI-Challenge/baseline2025>

## References

- [1] Minja Axelsson, Micol Spitale, and Hatice Gunes. 2024. "Oh, Sorry, I Think I Interrupted You": Designing Repair Strategies for Robotic Longitudinal Well-being Coaching. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*. 13–22.
- [2] Tadas Baltrušaitis, Amir Zadeh, Yao Chong Lim, and Louis-Philippe Morency. 2018. OpenFace 2.0: Facial Behavior Analysis Toolkit. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*. 59–66. doi:10.1109/FG.2018.00019
- [3] Atef Ben-Youssef, Chloé Clavel, and Slim Essid. 2019. Early detection of user engagement breakdown in spontaneous human-humanoid interaction. *IEEE Transactions on Affective Computing* 12, 3 (2019), 776–787.
- [4] Alexandra Bremers, Alexandria Pabst, Maria Teresa Parreira, and Wendy Ju. 2023. Using social cues to recognize task failures for hri: A review of current research and future directions. *arXiv preprint arXiv:2301.11972* (2023).
- [5] Kay Henning Brodersen, Cheng Soon Ong, Klaas Enno Stephan, and Joachim M Buhmann. 2010. The balanced accuracy and its posterior distribution. In *2010 20th international conference on pattern recognition*. IEEE, 3121–3124.
- [6] Kate Candon, Jesse Chen, Yoony Kim, Nathan Tsoi, and Marynel Vázquez. 2023. Nonverbal Human Signals Can Help Autonomous Agents Infer Human Preferences for Their Behavior. (2023).
- [7] Shiye Cao and Chien-Ming Huang. 2022. Understanding user reliance on AI in assisted decision-making. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–23.
- [8] Shiye Cao, Jiwon Moon, Amama Mahmood, Victor Nikhil Antony, Ziang Xiao, Anqi Liu, and Chien-Ming Huang. 2025. Interruption Handling for Conversational Robots. *Proceedings of the 2025 Robotics: Science and Systems Conference* (2025).
- [9] Jennifer Carlson and Robin R Murphy. 2005. How UGVs physically fail in the field. *IEEE Transactions on robotics* 21, 3 (2005), 423–437.
- [10] Connor Esterwood and Lionel P Robert. 2021. Do you still trust me? human-robot trust repair strategies. In *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*. IEEE, 183–188.
- [11] Florian Eyben, Martin Wöllmer, and Björn Schuller. 2010. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM international conference on Multimedia*. 1459–1462.
- [12] Kathrin Gerling, Denise Hebesberger, Christian Dondrup, Tobias Körtner, and Marc Hanheide. 2016. Robot deployment in long-term care: Case study on using a mobile robot to support physiotherapy. *Zeitschrift für Gerontologie und Geriatrie* 49 (2016), 288.
- [13] Manuel Giuliani, Nicole Mirnig, Gerald Stollnberger, Susanne Stadler, Roland Buchner, and Manfred Tscheligi. 2015. Systematic analysis of video data from different human-robot interaction studies: a categorization of social signals during error situations. *Frontiers in psychology* 6 (2015), 931.
- [14] Shanee Honig and Tal Oron-Gilad. 2018. Understanding and resolving failures in human-robot interaction: Literature review and model development. *Frontiers in psychology* 9 (2018), 861.
- [15] Chien-Ming Huang and Bilge Mutlu. 2013. The repertoire of robot behavior: Enabling robots to achieve interaction goals through social behavior. *Journal of Human-Robot Interaction* 2, 2 (2013), 80–102.
- [16] Ruben Janssens, Eva Verhelst, and Mathieu De Coster. 2024. Predicting Errors and Failures in Human-Robot Interaction from Multi-Modal Temporal Data. In *Proceedings of the 26th International Conference on Multimodal Interaction*. 671–676.
- [17] Ulas Berk Karli, Shiye Cao, and Chien-Ming Huang. 2023. "What If It Is Wrong": Effects of Power Dynamics and Trust Repair Strategy on Trust and Compliance in HRI. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. 271–280.
- [18] Dimosthenis Kontogiorgos, Minh Tran, Joakim Gustafson, and Mohammad Soleymani. 2021. A systematic cross-corpus analysis of human reactions to robot conversational failures. In *Proceedings of the 2021 International Conference on Multimodal Interaction*. 112–120.
- [19] Amama Mahmood, Shiye Cao, Maia Stiber, Victor Nikhil Antony, and Chien-Ming Huang. 2025. Voice assistants for health self-management: Designing for and with older adults. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–22.
- [20] Amama Mahmood, Junxiang Wang, and Chien-Ming Huang. 2024. Situated Understanding of Older Adults' Interactions with Voice Assistants: A Month-long In-home Study. *arXiv preprint arXiv:2403.02421* (2024).
- [21] Amama Mahmood, Junxiang Wang, Bingsheng Yao, Dakuo Wang, and Chien-Ming Huang. 2025. User interaction patterns and breakdowns in conversing with LLM-powered voice assistants. *International Journal of Human-Computer Studies* 195 (2025), 103406.
- [22] Nicole Mirnig, Manuel Giuliani, Gerald Stollnberger, Susanne Stadler, Roland Buchner, and Manfred Tscheligi. 2015. Impact of robot actions on social signals and reaction times in HRI error situations. In *Social Robotics: 7th International Conference, ICSR 2015, Paris, France, October 26-30, 2015, Proceedings 7*. Springer, 461–471.
- [23] David Portugal, Paulo Alvito, Eleni Christodoulou, George Samaras, and Jorge Dias. 2019. A study on the deployment of a service robot in an elderly care center. *International Journal of Social Robotics* 11 (2019), 317–341.
- [24] Pradip Pramanick and Silvia Rossi. 2024. PRISCA at ERR@ HRI 2024: Multimodal Representation Learning for Detecting Interaction Ruptures in HRI. In *Proceedings of the 26th International Conference on Multimodal Interaction*. 666–670.
- [25] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. 8748–8763.
- [26] Allison Sauppé and Bilge Mutlu. 2014. How social cues shape task coordination and communication. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. 97–108.
- [27] Brian Scassellati, Laura Boccanfuso, Chien-Ming Huang, Marilena Mademtz, Meiying Qin, Nicole Salomons, Pamela Ventola, and Frederick Shic. 2018. Improving social skills in children with ASD using a long-term, in-home social robot. *Science Robotics* 3, 21 (2018), eaat7544.
- [28] Lisa Scherf, Lisa Alina Gasche, Eya Chemangui, and Dorothea Koert. 2024. Are You Sure? Multi-Modal Human Decision Uncertainty Detection in Human-Robot Interaction. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*. 621–629.
- [29] Micol Spitale, Minja Axelsson, Neval Kara, and Hatice Gunes. 2023. Longitudinal Evolution of Coachees' Behavioural Responses to Interaction Ruptures in Robotic Positive Psychology Coaching. In *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 315–322.
- [30] Micol Spitale, Maria Teresa Parreira, Maia Stiber, Minja Axelsson, Neval Kara, Garima Kankariya, Chien-Ming Huang, Malte Jung, Wendy Ju, and Hatice Gunes. 2024. ERR@ HRI 2024 Challenge: Multimodal Detection of Errors and Failures in Human-Robot Interactions. In *Proceedings of the 26th International Conference on Multimodal Interaction*. 652–656.
- [31] Maia Stiber, Dan Bohus, and Sean Andrist. 2024. "Uh, This One?": Leveraging Behavioral Signals for Detecting Confusion during Physical Tasks. In *Proceedings of the 26th International Conference on Multimodal Interaction*. 194–203.
- [32] Maia Stiber, Russell Taylor, and Chien-Ming Huang. 2022. Modeling human response to robot errors for timely error detection. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 676–683.
- [33] Maia Stiber, Russell Taylor, and Chien-Ming Huang. 2025. Robot Error Awareness Through Human Reactions: Implementation, Evaluation, and Recommendations. *arXiv preprint arXiv:2501.05723* (2025).
- [34] Maia Stiber, Russell H Taylor, and Chien-Ming Huang. 2023. On using social signals to enable flexible error-aware hri. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. 222–230.
- [35] Aaqib Tabrez, Matthew B Luebbers, and Bradley Hayes. 2020. A survey of mental modeling techniques in human-robot teaming. *Current Robotics Reports* 1 (2020), 259–267.
- [36] Alessandro Vinciarelli, Maja Pantic, Hervé Bourlard, and Alex Pentland. 2008. Social signals, their function, and automatic analysis: a survey. In *Proceedings of the 10th international conference on Multimodal interfaces*. 61–68.
- [37] Alessandro Vinciarelli, Maja Pantic, Dirk Heylen, Catherine Pelachaud, Isabella Poggi, Francesca D'Errico, and Marc Schroeder. 2011. Bridging the gap between social animal and unsocial machine: A survey of social signal processing. *IEEE Transactions on Affective Computing* 3, 1 (2011), 69–87.
- [38] Lennart Wachowiak, Peter Tisnikar, Andrew Coles, Gerard Canal, and Oya Celiktutan. 2024. A Time Series Classification Pipeline for Detecting Interaction Ruptures in HRI Based on User Reactions. In *Proceedings of the 26th International Conference on Multimodal Interaction*. 657–665.
- [39] Jason R Wilson, Phyo Thuta Aung, and Isabelle Boucher. 2022. When to help? A multimodal architecture for recognizing when a user needs help from a social robot. In *International Conference on Social Robotics*. Springer, 253–266.