Don't Go Breaking My Trust: A Taxonomy of Interaction Failures for Transparent Robots

Beatriz Nogueira, João Ferreira, Hugo Simão, Filipa Rocha, Isabel Neto, João Guerreiro, Tiago Guerreiro
beatriz.c.nogueira@outlook.pt,jpferreira@lasige.di.fc.ul.pt,hugalexsimon@gmail.com
fgrocha@fc.ul.pt,aineto@fc.ul.pt
jpguerreiro@fc.ul.pt,tjguerreiro@ciencias.ulisboa.pt
LASIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal

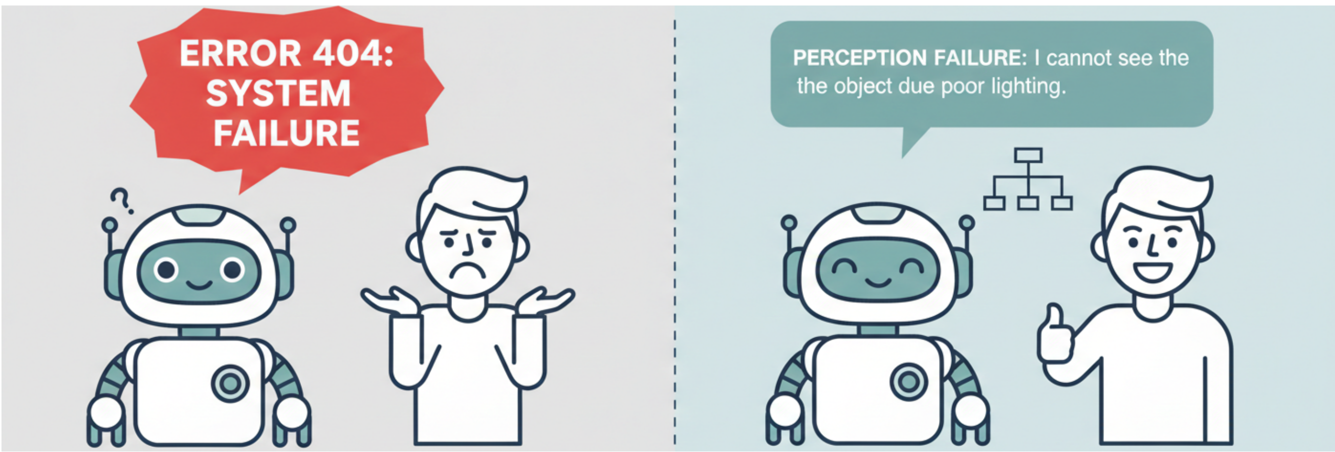

Figure 1: Shifting from technical "silent" failures to transparent, understandable communication. On the left, current HRI systems often provide generic error alerts that leave users confused and erode trust. On the right, our proposed multidimensional taxonomy allows robots to diagnose and communicate the specific nature of a failure (e.g., a Perception Failure), enabling context-aware explanations that sustain safe and effective human-robot collaboration.

Abstract

As robots become integrated into everyday environments, interaction failures are inevitable. However, current robotics design often treats failures as technical bugs rather than interactive mistakes. When these errors are presented in technical terms, users might not be able to understand them, reducing trust and the ease of the interaction. To address this, we present a preliminary, multidimensional taxonomy of HRI failures. Based on team brainstorming and literature, our framework classifies failures across three core system capabilities: Perception, Cognition, and Execution. Crucially, it attributes these failures not solely to the robot, but also to human and environmental factors, while assessing their impact on user trust and safety. By pinpointing the exact nature and origin of an error, this taxonomy provides the vocabulary needed to equip robots with targeted, context-aware explanations, advancing the design of transparent and understandable human-robot interaction.

CCS Concepts

• **Human-centered computing** → *Interaction paradigms.*

Keywords

HRI Taxonomy, Trust, Transparency, Failures

1 Introduction

Regardless of the technical and social efforts to make robots reliable, the reality in human-robot interaction (HRI) is that they continue to fail and frequently encounter unexpected errors and misinterpretations [16, 26]. Crucially, every technical error carries a consequence for the interaction that often goes unmeasured: the perceived dependability. As we progress toward a more robot-integrated environment, with robots moving from controlled laboratory settings to real-world contexts [8, 21, 23, 34], the importance of exploring and managing these vulnerabilities increases.

A failure in this context is broadly defined as a robot behavior that deviates from the user's desired or expected outcome [30]. Given the complexity of dynamic everyday environments [16], such failures are inherent to interaction and can be triggered by the robot's system [15, 19], environmental anomalies [9, 29], or human actions [14, 32]. Since eliminating these failures across different contexts is unlikely to be feasible, the focus of modern HRI should shift to how robots manage, mitigate, and communicate them.

Current robotics designs often treat failure simply as a system error to be eliminated [2–4, 25]. However, for a user, a failure can mean more than that, possibly further undermining their already limited trust in the robotic agent [21]. To restore this trust and ensure continued safety, the user must understand how and why the failure occurred, and how to recover from it [6, 10, 33]. By designing robots capable of transparently communicating their intent, reasoning processes, and system limitations, we can transform these failures from interaction- and trust-ending errors into regular, understandable events.

To better understand failures, prior work has sought to define them, with a focus on user-centered [16], safety [18], and mobile robots [5]. However, a comprehensive, structured taxonomy encompassing various failure types, their consequences for factors such as safety and trust, and their recoverability in HRI remains underexplored. These dimensions are particularly relevant during live interactions, where an unexplained error does more than affect a task; it fundamentally alters the user's perception and confidence towards the robot. Without a structured way to classify these failures, they remain 'silent,' leading to a decline in trust that might be difficult to repair.

This work aims to address this gap by developing a systematic and comprehensive failure taxonomy that integrates multiple perspectives. By providing a structured framework for understanding and categorizing failures, this taxonomy can help researchers and practitioners identify, discuss, quantify, compare systems and their interactive behaviors, and ultimately mitigate failures in human-robot interaction while promoting user trust.

2 Methodology

In this section, we describe the (ongoing) process to develop and validate a comprehensive and robust failure taxonomy for HRI. To address both the technical and social dimensions of HRI, we are following a five-phase iterative approach. Currently, the first two phases are complete, yielding the preliminary taxonomy presented in this paper, while the third phase is ongoing.

2.1 Initial Review of Existing Taxonomies (Completed)

To establish a foundation for our taxonomy, we conducted a comprehensive review of existing taxonomies related to failures in robotics, assistive robots, service robots, and HRI. We used the Google Scholar search engine to find relevant articles. To achieve this, we used specific key words which included "HRI Taxonomy", "robotics taxonomy", "Failure taxonomy", etc. This process allowed us to understand the various dimensions in which failures have historically been categorized and informed our initial framework for structuring failure types. We documented and organized these insights using a Miro board, which served as a collaborative space for mapping and refining our taxonomy.

2.2 Preliminary Taxonomy Development (Completed)

Building on the initial review, we initiated the taxonomy development phase by synthesizing the extracted failure categories into a preliminary framework. This early version of the taxonomy included broad failure types based on previously identified patterns. We then iteratively refined them to better fit the HRI context.

To ensure a multidisciplinary perspective, we engaged an interdisciplinary team of 7 researchers with expertise in developing conceptual frameworks, HRI, Robotics, HCI, Accessibility, Design, Virtual and Augmented Reality, and Computer Vision. The team contributed equally in the discussion process.

Utilizing a collaborative digital workspace (Miro), our team mapped and iteratively refined the categories through multiple rounds of discussion to better align with the specific nuances of HRI. Each iteration aimed to improve clarity, coverage, and applicability of the taxonomy. Adjustments were made in response to feedback from these sessions. Categories were merged, subdivided, or redefined to ensure that the taxonomy was both comprehensive and applicable across different HRI contexts.

2.3 Scoping Review (Ongoing)

To ensure our framework objectively captures the full spectrum of documented real-world HRI failures, we are currently conducting a scoping review that grounds our taxonomy in a comprehensive empirical research base. This review follows the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) statement [24, 31]to ensure transparency and reproducibility.

We queried three major academic databases: Science Direct, IEEE Xplore, and ACM Digital Library, using the following search string:

```
(robot*) AND (interaction* OR Collaboration*
OR user* OR experience OR cooperation* OR
assistive* OR rehabilitation* OR service
OR social*) AND (fail* OR malfunction OR
error OR breakdown OR incident OR violation)
```

To define the scope of our review, we applied the inclusion and exclusion criteria detailed in Table 1. We prioritized studies involving physical robots to ensure the taxonomy reflects the challenges of real-world deployment.

The review is guided by the following Research Questions (RQs): RQ1: What types of failures have been reported in human-robot interaction, and how are they described or classified across studies? RQ2: What types of failures are under-represented or overlooked in the literature? RQ3: What factors contribute to these failures (e.g., robot, human, and environmental aspects)?

To systematically analyze the included papers, we developed a data extraction spreadsheet. To strengthen the validity of our findings, each paper was independently annotated by two of the four researchers involved in this phase. Any discrepancies or conflicts in the annotations were resolved through collaborative discussion with the full research team, ensuring a high level of inter-annotator agreement. For each selected study, we extract contextual metadata, including the study type, study context (e.g., healthcare, public space), robot model, target users, and the robot's level of autonomy during the interaction, as well as the number of agents in the interaction (humans and/or robots). Crucially, to directly answer our Research Questions and map findings to our preliminary taxonomy, we extract specific failure-related metrics for every documented incident. This metrics include: How the authors name and define the failure; When the failure occurred during the interaction, and whether the primary agent at fault was the human, the robot, or the environment; The reported effects of the failure on the interaction, specifically regarding safety, usability, and user trust, which has been shown to be significantly impacted by both technical and social robot errors [11, 12, 19, 20]; and if there were any recovery actions taken by the system, such as providing an explanation,

Table 1: Inclusion and Exclusion Criteria for the Scoping Review

Category	Inclusion Criteria	Exclusion Criteria
Population	Humans interacting with **physical** robots (social, service, assistive).	Simulated robots, virtual agents, or pure simulation environments.
Concept	Explicit focus on failures, errors, or unexpected behaviors.	Failures mentioned only anecdotally; pure algorithm validation without interaction data.
Context	Failures occurring during active human-robot interaction (or in shared environments).	Technical failures analyzed in isolation (e.g., benchmarking) without human impact.
Study Type	Empirical studies, case studies, experiments, and peer-reviewed theoretical frameworks.	Pre-prints, editorials, opinion pieces, challenges/protocols.
Publication	Peer-reviewed (Journal/Conference); English; Published $\geq$ 2015.	Non-English; Unavailable full text; Gray literature.

attempting a system reset, or prompting a human takeover. The purpose of this ongoing review is to systematically map existing evidence, identify underrepresented failure types, and extract specific use cases to test against our framework.

2.4 Taxonomy Iteration (Future Work)

Once the scoping review is complete, the extracted data will be used to expand, validate, and improve our preliminary taxonomy. Categories will be adjusted, merged, or subdivided based on the newly mapped evidence to ensure comprehensive coverage of technical, social, and environmental failure aspects.

2.5 Validation and Evaluation (Future Work)

Lastly, to evaluate the robustness and applicability of the taxonomy, we will conduct a validation phase in which we systematically categorize documented robotic and HRI failures using it. This involved reviewing failure reports from academic literature, case studies, and real-world deployments. We will also conduct focus groups with roboticists, designers, engineers, and HRI specialists. They will be presented with a series of HRI failure scenarios and asked to categorize them using our framework to assess its feasibility and perceived clarity. This iterative validation process will ensure that the taxonomy effectively captures real-world failures while maintaining clarity and consistency. As a result, our final taxonomy provides a structured and comprehensive framework for categorizing failures in HRI, supporting both research and practical applications in failure analysis and mitigation.

3 Preliminary Taxonomy

Our preliminary taxonomy classifies failures in HRI based on the moment they occur and their underlying cause (Fig. 2).

The primary dimensions of our framework: Perception, Cognition, and Execution, are grounded in the PCA (Perception-Cognition-Action) model [1, 13] commonly used in human factors engineering and cognitive psychology, and often related to the Sense-Plan-Act cycle in robotics [27]. This model provides a robust alternative to purely social error taxonomies by focusing on the underlying functional process of the robotic system. While existing work often categorizes errors by their social outcome (e.g., a "norm violation"), our taxonomy identifies the technical root (e.g., a "perception failure" of a social cue), which is essential for a robot to generate an accurate, transparent explanation for the user.

A core premise of our framework is that interaction failures are not solely the robot's fault. They can be triggered by three distinct agents: the Robot, the Human, and the Environment, with the possibility that more than one agent can initiate the failure. In a case where a failure leads to another, we apply a root cause principle, where a failure is attributed to the earliest stage in the PCA cycle.

Failures are grouped into three primary dimensions corresponding to the fundamental capabilities of a robotic (or human) system: Perception, Cognition, and Execution. The three dimensions are defined as follows:

- **Perception Failures**: These occur when a robot or user misinterprets sensory information. This dimension includes actors' perception errors (e.g., the robot misidentifies a human user or vice versa) and environmental perception errors (e.g., the robot incorrectly interprets its surroundings and fails to adapt, or sunlight affects the robot's sensors). Common triggers include sensor limitations, misperceptions of the robot's capabilities, algorithmic miscalculations, or external environmental factors such as poor lighting and excessive noise.
- **Execution Failures**: These occur during the attempt to perform an action. This includes mechanical failures (e.g., malfunctioning motors or actuators), coordination failures (e.g., failing to synchronize movements with a user, resulting in failed handovers), and communication failures (e.g., providing incorrect or inadequate feedback; the user using the wrong or no input, etc.). These errors often stem from hardware deterioration, real-time movement coordination errors, or mishaps in verbal and non-verbal signaling.
- **Cognition Failures**: These stem from incorrect processing, reasoning, or memory retention, which impedes the robot's

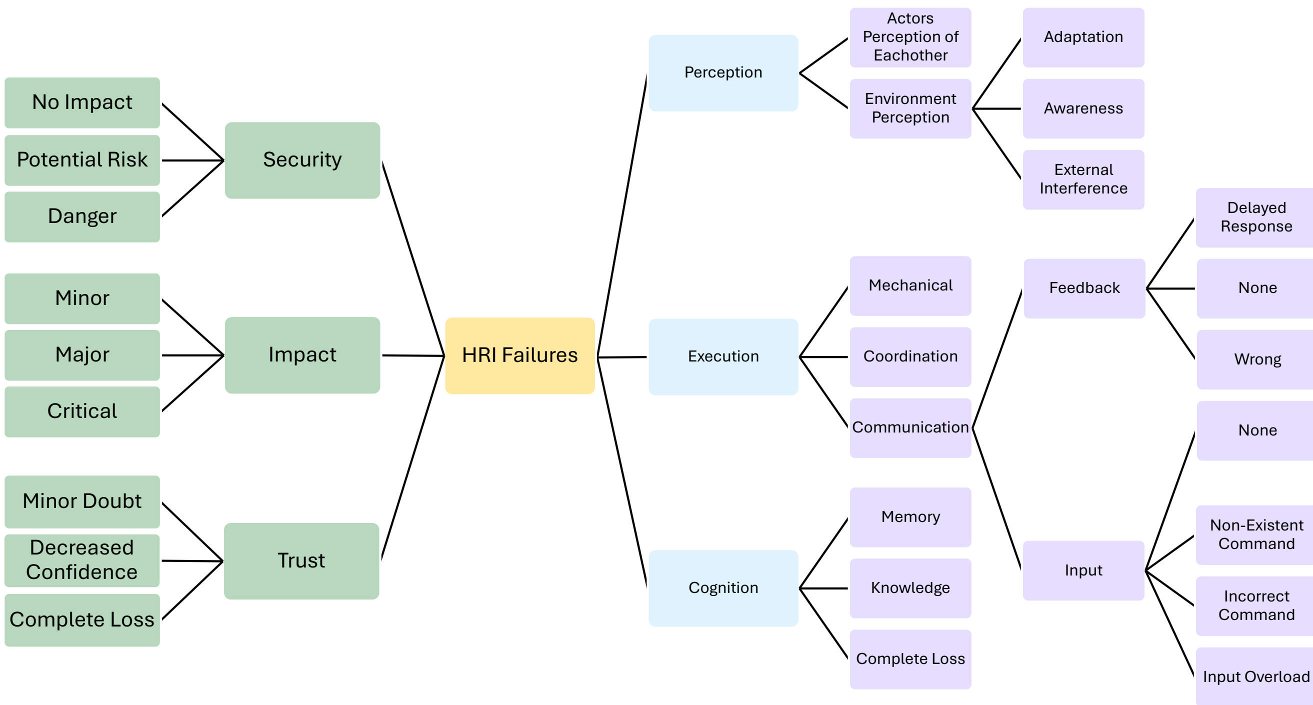

Figure 2: The proposed multidimensional HRI Failure Taxonomy. The framework maps failures across functional dimensions (Perception, Execution, Cognition) and evaluates their downstream consequences across three critical impact axes: Security, Impact, and User Trust.

ability to understand and adapt to human interaction. Subcategories include adaptation failures (failing to respond to specific user needs), memory failures (forgetting previously learned information), and knowledge gaps (lacking the necessary information to make appropriate decisions; "forgetting" information, etc).

Beyond categorizing the type of failure, our taxonomy integrates attributes to assess the failure's overall significance on the interaction. We argue that these attributes are not standalone, but rather emergent properties that influence one another and dictate the necessary recovery strategy. A failure can have repercussions for users' safety, ranging from no impact to those that introduce potential risks to those that create immediate danger. Connected to this is the effect on users' trust, which serves as a primary proxy for the user's perception of the robot. This assesses the failure's effect on user reliance, evaluating whether it causes minor user doubt, decreased confidence, or a complete loss of trust. Additionally, we have the overall Impact, which encompasses these categories, as well as error mitigation and recovery. It can range from minor to major to critical, depending on how low or severe these aspects are.

Classifying HRI failures through this multidimensional lens is a necessary prerequisite for designing transparent, understandable robots, as a system can only provide meaningful explanations and restore user trust when it accurately understands the specific nature and origin of its own failure.

4 Discussion and Future Work

While our scoping review is still ongoing, our initial full-text screening has already revealed a significant technocentric bias in how the HRI community currently reports failures. We have identified that the vast majority of research articles treat failures and errors as purely quantitative, technical metrics. Frequently, failures are mentioned only in passing to benchmark algorithmic performance, with one system deemed superior to another solely based on a lower error percentage [17, 28]. These technical validations routinely omit crucial contextual details regarding how the failure actually unfolded during the human-robot interaction. Consequently, there is a distinct lack of qualitative analysis concerning how these failures are perceived, experienced, and interpreted by the users themselves.

4.1 The Need for a Transparency-Driven Taxonomy

This preliminary finding underscores a critical gap in the field: treating a failure merely as a system bug or an algorithmic error rate ignores the inherent social contract of HRI. If a robot's internal logic treats an error as a statistical anomaly, it cannot effectively communicate to the user what went wrong [7, 22].

To design transparent and understandable robots, the field must shift from analyzing failures purely as technical defects to understanding them as critical interactive events. Our proposed taxonomy facilitates this shift. By classifying failures not just by their technical source (e.g., sensor noise), but also by their interactive dimension (Perception, Execution, Cognition) and their impact on the user (Trust, Safety), we provide designers with the specific vocabulary needed to trigger appropriate mitigation strategies. Furthermore, by explicitly expanding the taxonomy to encompass human and environmental triggers alongside robot errors, the framework facilitates the anticipation and prediction of interaction failures. Recognizing these external sources enables designers to proactively predict when and how failures might occur, enabling more robust, timely mitigation. For instance, an Execution Failure (e.g., dropping an object) requires a fundamentally different explanation and recovery strategy than a Cognition Failure (e.g., misinterpreting a user's intent or forgetting a preference).

4.2 Bridging Taxonomy and Transparent Communication

Moving forward, our immediate work involves completing the full scoping review to thoroughly validate and populate the proposed taxonomy. Beyond this, our goal is to leverage the taxonomy as a foundational framework for designing active, transparent error-recovery behaviors.

As robots become increasingly autonomous in everyday environments, it is no longer enough for them to fail silently; they must be understandable in their failures. A generic error alert, or a complete lack of communication, leaves users confused about what went wrong and does little to repair trust. By pinpointing exactly where the failure occurred (Perception, Cognition, or Execution) and what triggered it (Robot, Human, or Environment), designers can equip robots with targeted, context-aware explanations and mitigation strategies. This framework moves failure handling away from simple technical debugging and toward proactive, understandable communication, ensuring that when robots inevitably fail, they can explain why in a way that preserves safe and effective human-robot collaboration.

Moving forward, our goal is to leverage this taxonomy not merely as an analytical tool, but as a foundational mechanism for active failure prediction and recovery. Because a mechanical execution error requires a fundamentally different response than a cognitive misunderstanding of user intent, it is clear that each taxonomy category demands a distinct mitigation strategy. We plan to explore how robots can use this framework to self-diagnose the nature of an error and dynamically adapt their communication. Future studies will focus on linking these specific failure categories to tailored explanations of the robot's reasoning processes and underlying uncertainties. Ultimately, by accurately mapping how and why failures occur, we aim to empower systems to anticipate errors, recover gracefully, and sustain user trust through transparent interaction.

Acknowledgments

This work is partially funded by FCT, through scholarships 2025.03576.BD and 2025.02361.BD, the LASIGE Research Unit, ref. UID/00408/2025, DOI https://doi.org/10.54499/UID/00408/2025, and through the European Union's Horizon Europe programme project "HARIA: Human-Robot Sensorimotor Augmentation - Wearable Sensorimotor Interfaces And Supernumerary Robotic Limbs For Humans With Upper-Limb Disabilities", under grant agreement No. 101070292.

References

[1] [n. d.]. Introducing the PCA Model in Human Factors Engineering - Research Collective. https://research-collective.com/introducing-the-pca-model-in-human-factors-engineering/

[2] Cristhian A. Aguilera-Carrasco, Luis Felipe Gonzalez-Bohme, Francisco Valdes, Francisco Javier Quitral-Zapata, and Bogdan Raducanu. 2023. A Hand-Drawn Language for Human-Robot Collaboration in Wood Stereotomy. *IEEE Access* 11 (2023), 100975–100985. doi:10.1109/ACCESS.2023.3314337

[3] Martin K. Burns, DIngyi Pei, and Ramana Vinjamuri. 2019. Dynamic Control of Virtual Hand Grasp Using Spatiotemporal Synergies. *IEEE Access* 7 (2019), 112327–112338. doi:10.1109/ACCESS.2019.2932956

[4] Jennifer Carlson and Robin R. Murphy. 2005. How UGVs physically fail in the field. *IEEE Transactions on Robotics* 21, 3 (6 2005), 423–437. doi:10.1109/TRO.2004.838027

[5] Jennifer Carlson, Robin R. Murphy, and Andrew Nelson. 2004. Follow-up analysis of mobile robot failures. *Proceedings - IEEE International Conference on Robotics and Automation* 2004, 5 (2004), 4987–4994. doi:10.1109/robot.2004.1302508

[6] Devleena Das, Siddhartha Banerjee, and Sonia Chernova. 2021. Explainable AI for robot failures: Generating explanations that improve user assistance in fault recovery. *ACM/IEEE International Conference on Human-Robot Interaction* (3 2021), 351–360. doi:10.1145/3434073.3444657

[7] Maartje M A De Graaf and Bertram F Malle. 2017. How People Explain Action (and Autonomous Intelligent Systems Should Too).. In *AAAI Fall Symposia*, Vol. 8.

[8] Ha Manh Do, Minh Pham, Weihua Sheng, Dan Yang, and Meiqin Liu. 2018. RiSH: A robot-integrated smart home for elderly care. *Robotics and Autonomous Systems* 101, 1-2 (3 2018), 74–92. doi:10.1016/j.robot.2017.12.008

[9] Matthias Eder, Bettina Kubicek, and Gerald Steinbauer-Wagner. 2026. Effect of efficiently computed explanations for robot motion planning failures in human–robot interaction. *Robotics and Autonomous Systems* 195, 4 (1 2026), 105224. doi:10.1016/j.robot.2025.105224

[10] Connor Esterwood and Lionel P. Robert. 2022. A Literature Review of Trust Repair in HRI. *RO-MAN 2022 - 31st IEEE International Conference on Robot and Human Interactive Communication: Social, Asocial, and Antisocial Robots* (2022), 1641–1646. doi:10.1109/RO-MAN53752.2022.9900667

[11] Ioanna Giorgi, Francesca Ausilia Tirotto, Oksana Hagen, Farida Aider, Mario Gianni, Marco Palomino, and Giovanni L. Masala. 2022. Friendly But Faulty: A Pilot Study on the Perceived Trust of Older Adults in a Social Robot. *IEEE Access* 10 (2022), 92084–92096. doi:10.1109/ACCESS.2022.3202942

[12] Hideki Garcia Goo, Bob R. Schadenberg, Jan Kolkmeier, Khiet P. Truong, and Vanessa Evers. 2025. Being Sorry is the Hardest Thing: How Robots can Apologize and Learn from Mistakes to Restore People's Trust. *Conference on Human Factors in Computing Systems - Proceedings* (4 2025). doi:10.1145/3706599.3719739

[13] Pascal Haazebroek, Saskia Van Dantzig, and Bernhard Hommel. 2011. A computational model of perception and action for cognitive robotics. *Cognitive Processing 2011 12:4* 12, 4 (5 2011), 355–365. doi:10.1007/s10339-011-0408-x

[14] Andreas Kornmaaler Hansen, Valeria Villani, Andrea Pupa, and Astrid Heidemann Lassen. 2025. Introducing Novice Operators to Collaborative Robots: A Hands-On Approach for Learning and Training. *IEEE Transactions on Automation Science and Engineering* 22 (2025), 3933–3946. doi:10.1109/TASE.2024.3403709

[15] Daniel Herrera, Flavio Roberti, Marcos Toibero, and Ricardo Carelli. 2017. Human interaction dynamics for its use in mobile robotics: Impedance control for leader-follower formation. *IEEE/CAA Journal of Automatica Sinica* 4, 4 (10 2017), 696–703. doi:10.1109/JAS.2017.7510631

[16] Shanee Honig and Tal Oron-Gilad. 2018. Understanding and resolving failures in human-robot interaction: Literature review and model development. *Frontiers in Psychology* 9, JUN (6 2018), 351644. doi:10.3389/fpsyg.2018.00861

[17] Haiyi Kong, Chenguang Yang, Guang Li, and Shi Lu Dai. 2020. A sEMG-Based Shared Control System with No-Target Obstacle Avoidance for Omnidirectional Mobile Robots. *IEEE Access* 8 (2020), 26030–26040. doi:10.1109/ACCESS.2020.2970468

[18] Brian H.W. GUO;Yonger ZUO;Yang Miang Goh;Jae-Yong Lim;. 2024. Errors in Human-Robot Interaction Accidents: A Taxonomy and Network Analysis. *International conference on construction engineering and project management* (2024), 1088–1095. doi:10.6106/ICCEPM.2024.1088

[19] Joseph B. Lyons, Izz aldin Hamdan, and Thy Q. Vo. 2023. Explanations and trust: What happens to trust when a robot partner does something unexpected? *Computers in Human Behavior* 138, 7729 (1 2023), 107473. doi:10.1016/j.chb.2022.107473

[20] Atsuya Matsumoto, Takashige Suzuki, Chi Lan Yang, Takuji Narumi, and Hideaki Kuzuoka. 2025. In Touch We Decide: Physical Touch by Embodied Virtual Agent Increases the Acceptability of Advice. *IEEE Transactions on Visualization and Computer Graphics* 31, 5 (2025), 2839–2847. doi:10.1109/TVCG.2025.3549559

[21] Linda Miller, Johannes Kraus, Franziska Babel, and Martin Baumann. 2021. More Than a Feeling—Interrelation of Trust Layers in Human-Robot Interaction and the Role of User Dispositions and State Anxiety. *Frontiers in Psychology* 12 (4 2021), 592711. doi:10.3389/fpsyg.2021.592711

[22] Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* 267, 2 (2 2019), 1–38. doi:10.1016/j.artint.2018.07.007

[23] Beatriz Nogueira, Hugo Simão, Alexandre Bernardino, Jodi Forlizzi, and Tiago Guerreiro. 2025. Putting the Age in Agency A Case Study of an Older Adult Controlling a Telepresence Robot in a Family Setting. *DIS 2025 - Companion Proceedings of the 2025 ACM Designing Interactive Systems Conference: Designing for a Sustainable Ocean* (7 2025), 508–512. doi:10.1145/3715668.3736344

[24] Matthew J Page, Joanne E Mckenzie, Patrick M Bossuyt, Isabelle Boutron, Tammy C Hoffmann, Cynthia D Mulrow, Larissa Shamseer, Jennifer M Tetzlaff, Elie A Akl, Sue E Brennan, Roger Chou, Julie Glanville, Jeremy M Grimshaw, Asbjørn Hróbjartsson, Manoj M Lalu, Tianjing Li, Elizabeth W Loder, Evan Mayo-Wilson, Steve Mcdonald, Luke A Mcguinness, Lesley A Stewart, James Thomas, Andrea C Tricco, Vivian A Welch, Penny Whiting, and David Moher. [n. d.]. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. ([n. d.]). doi:10.1136/bmj.n71

[25] Eric Price, Guilherme Lawless, Roman Ludwig, Igor Martinovic, Heinrich H. Bulthoff, Michael J. Black, and Aamir Ahmad. 2018. Deep Neural Network-Based Cooperative Visual Tracking Through Multiple Micro Aerial Vehicles. *IEEE Robotics and Automation Letters* 3, 4 (10 2018), 3193–3200. doi:10.1109/LRA.2018.2850224

[26] Micol Spitale, San Jose, Costa Rica, Maria Teresa Parreira Maia Stiber, Minja Axelsson Neval Kara, Garima Kankariya, Chien-Ming Huang Malte Jung Wendy Ju, Hatice Gunes, Maria Teresa Parreira, Maia Stiber, Minja Axelsson, Neval Kara, Chien-Ming Huang, Malte Jung, and Wendy Ju. 2024. ERR@HRI 2024 Challenge: Multimodal Detection of Errors and Failures in Human-Robot Interactions Citation in BibTeX format ICMI '24: INTERNATIONAL CONFERENCE ON MULTIMODAL INTERACTION Challenge: Multimodal Detection of Errors and Failures in Human-Robot I. . . . (2024). doi:10.1145/3678957.3689030

[27] Ankit Srivastava. 2019. sense. (2019). https://www.researchgate.net/profile/Ankit-Srivastava-46/publication/349248621_Sense-Plan-Act_in_Robotic_Applications/links/60268a00a6fdcc37a81e454c/Sense-Plan-Act-in-Robotic-Applications.pdf

[28] Ning Sun, Menglin Cao, Yuhao Chen, Yanni Chen, Jue Wang, Qiong Wang, Xiaocong Chen, and Tian Liu. 2022. Continuous Estimation of Human Knee Joint Angles by Fusing Kinematic and Myoelectric Signals. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 30 (2022), 2446–2455. doi:10.1109/TNSRE.2022.3200485

[29] Krzysztof Adam Szczurek, Raul Marin Prades, Eloise Matheson, Jose Rodriguez-Nogueira, and Mario Di Castro. 2022. Mixed Reality Human-Robot Interface With Adaptive Communications Congestion Control for the Teleoperation of Mobile Redundant Manipulators in Hazardous Environments. *IEEE Access* 10 (2022), 87182–87216. doi:10.1109/ACCESS.2022.3198984

[30] Leimin Tian and Sharon Oviatt. 2021. ACM Reference format: Leimin Tian and Sharon Oviatt. 2021. A Taxonomy of Social Errors in Human-Robot Interaction. *ACM Trans. Hum.-Robot Interact* 10, 2 (2021). doi:10.1145/3439720

[31] Andrea C. Tricco, Erin Lillie, Wasifa Zarin, Kelly K. O'Brien, Heather Colquhoun, Danielle Levac, David Moher, Micah D.J. Peters, Tanya Horsley, Laura Weeks, Susanne Hempel, Elie A. Akl, Christine Chang, Jessie McGowan, Lesley Stewart, Lisa Hartling, Adrian Aldcroft, Michael G. Wilson, Chantelle Garritty, Simon Lewin, Christina M. Godfrey, Marilyn T. MacDonald, Etienne V. Langlois, Karla Soares-Weiser, Jo Moriarty, Tammy Clifford, Özge Tunçalp, and Sharon E. Straus. 2018. PRISMA Extension for Scoping Reviews (PRISMA-ScR): Checklist and Explanation. *Annals of internal medicine* 169, 7 (10 2018), 467–473. doi:10.7326/M18-0850

[32] Angelika Warmbein, Lucas Hübner, Ivanka Rathgeber, Amrei Christin Mehler-Klamt, Jana Huber, Ines Schroeder, Christina Scharf, Marcus Gutmann, Johanna Biebl, Kirsi Manz, Eduard Kraft, Inge Eberl, Michael Zoller, and Uli Fischer. 2024. Robot-assisted early mobilization for intensive care unit patients: Feasibility and first-time clinical use. *International Journal of Nursing Studies* 152 (4 2024), 104702. doi:10.1016/j.ijnurstu.2024.104702

[33] Elmira Yadollahi, Fethiye Irmak Dogan, Yujing Zhang, Beatriz Nogueira, Tiago Guerreiro, Shelly Levy Tzedek, and Iolanda Leite. 2025. Expectations, Explanations, and Embodiment: Attempts at Robot Failure Recovery. (4 2025). http://arxiv.org/abs/2504.07266

[34] Xinyan Yu, Yiyuan Wang, Tram Thi Minh Tran, Yi Zhao, Julie Stephany Berrio Perez, Marius Hoggenmüller, Justine Humphry, Lian Loke, Lynn Masuda, Callum Parker, Martin Tomitsch, and Stewart Worrall. 2023. Robots in the Wild: Contextually-Adaptive Human-Robot Interactions in Urban Public Environments. *ACM International Conference Proceeding Series* 23 (12 2023), 701–705. doi:10.1145/3638380.3638440

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009