Promoting Trust in Industrial Human-Robot Collaboration through Preference-Based Optimization

Campagna, Giulio; Lagomarsino, Marta; Lorenzini, Marta; Chrysostomou, Dimitrios; Rehm, Matthias; Ajoudani, Arash

Published in:
IEEE Robotics and Automation Letters

DOI (link to publication from Publisher):
10.1109/LRA.2024.3455792

Publication date:
2024

Document Version
Accepted author manuscript, peer reviewed version

Promoting Trust in Industrial Human-Robot Collaboration through Preference-Based Optimization

Giulio Campagna[1], Marta Lagomarsino[2], Marta Lorenzini[2], Dimitrios Chrysostomou[3], Matthias Rehm[1], Arash Ajoudani[2]

***Abstract*— This paper proposes a novel theoretical framework for promoting trust in human-robot collaboration (HRC). The framework exploits Preference-Based Optimization (PBO) and focuses on three key interaction parameters: robot velocity profile, human-robot separation distance, and vertical proximity to the user's head. By iteratively refining these parameters based on qualitative feedback from human collaborators, the system dynamically adapts robot trajectories. This personalization aims to enhance users' confidence in the robot's actions and foster a more trusting collaborative environment. In our user study with fourteen participants, we simulated a chemical industrial scenario for the HRC task. Results suggest that the framework effectively promotes human operator confidence in the robot assistant, particularly for individuals with limited prior experience in robotics.**

I. INTRODUCTION

As the manufacturing landscape evolves toward a new human-centric focus, the importance of trust in Human-Robot Collaboration (HRC) has become increasingly apparent. In the Industry 5.0 era, the focus shifts from mere automated production to creating a symbiotic and more harmonious relationship between operators and robotic counterparts [1]. A key aspect of this change is the concept of trust, a critical factor that shapes the dynamics of industrial HRC as it defines the level of effective collaboration in the workspace.

As defined in [2], trust is the expectation and confidence in the automation to perform its tasks appropriately, while [3] highlights the multifaceted nature of trust and its implications for system design and optimization. These early studies emphasized the necessity for the right balance of trust between humans and automated machines; while high levels of trust facilitate seamless collaboration, excessive or insufficient trust can lead to potential safety hazards or reduced efficiency [4].

Previous research has explored numerous perspectives for enhancing trust in HRC, including analyzing the impact of robot appearance [5], assessing its behavioral style [6], and implementing humanized dialogue protocols [7]. Despite the significance of trust during HRC, most prior research has primarily relied on post-hoc evaluations using questionnaires [8], [9], which offer a generic assessment of the interaction and trust dynamics without capturing real-time nuances. To effectively adapt their behavior in real-time, robots must be able to respond to fluctuations in trust as they occur during the collaboration.

A few approaches have been proposed to tailor the robot's functioning online based on the human psycho-physical state. In [10], human physical and cognitive factors and robot expenses are considered in a unified reinforcement-learning framework, which learns on-the-fly interaction parameters' values to align with user preferences and maximize the conjoint efficiency. In [11], HRC was treated as a repeated non-cooperative game with different objectives for the actors: stress minimization for the human and productivity maximization for the robot. The pace of the robot was then adjusted based on the estimated collaboration state. A novel human-in-the-loop framework that adjusts robot proximity and reactivity based on monitoring human attention and psycho-physical state is proposed in [12]. On-the-fly alterations of the robot path to enhance human comfort are combined with a multi-objective optimization to adapt the robot's trajectory execution time and smoothness based on the current stress level. However, assumptions about the users' tastes and inclinations during HRC were made in all these studies.

To empower individuals with choice, one promising approach involves using the *Preference-Based Optimization* (PBO) algorithms. These algorithms leverage human feedback to customize and adapt the behavior of robots according to individual preferences [13]. Given the subjective nature of perceiving robots during interactions, these algorithms can enhance the HRC experience by integrating human input into the optimization process, resulting in a more personalized and trust-inspiring interaction. The iterative nature of PBO allows for a gradual and guided convergence towards a preferred robot behavior, reducing the likelihood of sudden, unsettling changes that could compromise trust during HRC. Moreover, the optimization process itself can serve as a trust-building exercise, as the human operator witnesses the robot's performance improving over time following their own feedback. Several studies have explored the application of PBO algorithms in various HRC contexts. [14] utilized the C-GLISp algorithm, a global optimization approach that exploits pairwise comparisons to learn and adapt to human preferences. Similarly, [15] employed a PBO approach to optimize the robot's behavior in a collaborative sealing task, proving that trust and user satisfaction were improved compared to a non-adaptive baseline. Recent studies have also explored Preference-Based Learning approaches. In [16], a framework for personalized gait optimization in lower-body exoskeletons adapts walking trajectories based on user comfort preferences.

The work was supported by the ERC-StG Ergo-Lean (Grant No. 850932) and the Independent Research Fund Denmark (Grant No. 1032-00311B).

Corresponding author's email: gica@create.aau.dk

[1] G.Campagna and M.Rehm are with the Human-Robot Interaction Lab., Technical Faculty of IT and Design, Aalborg University, Denmark.

[2] M.Lagomarsino, M.Lorenzini, A.Ajoudani are with the Human-Robot Interfaces and Interaction Lab., Istituto Italiano di Tecnologia (IIT), Italy.

[3] D.Chrysostomou is with the Smart Production Lab., Faculty of Engineering and Natural Sciences, Aalborg University, Denmark.

In [17], a preference-based method learns reward functions using human feedback for robot trajectory comparisons. Despite the growing interest in PBO algorithms for trust-building in HRC, a key challenge about the application of PBO approaches to complex interactive human-robot tasks remains.

In this paper, we propose a novel framework that harnesses the power of PBO algorithms to enhance trust in HRC. The GLISp optimization algorithm is implemented to adapt the robot's trajectory to human preferences by iteratively refining 3 interaction parameters: robot velocity, separation distance, and vertical proximity to the user's head. These parameters were chosen for their validity across different scenarios and their known impact on trust during robot interactions [18], [19]. Studies show they influence human confidence in the robot's actions [20], [21], but also their combination can alter human perception and comfort [12]. Therefore, we design a system that does not optimize each interaction parameter in isolation but rather considers them collectively to identify the preferred parameters' combination. An analysis is conducted to assess the convergence of the optimization process and the alignment of the algorithm's outputs with human preferences. In addition, to determine how the parameter interdependencies affect the ability to select trust-promoting robot behaviors, we compared the combination of interaction parameters obtained with PBO and those selected in isolation. A benchmark scenario involves a robotic arm assisting a human in pouring chemicals, a process demanding precise coordination, situational awareness, and, most importantly, trust. An exhaustive subjective evaluation using standard and custom questionnaires assesses these approaches' effectiveness in capturing and adapting to individual trust requirements in HRC.

The main contributions of this work are twofold:

- Definition and optimization using the PBO algorithm of 3 key parameters that directly impact human-robot interaction and determine human trust level in robotic assistants;
- Adaptation of the robot's trajectory based on PBO-learned parameters to enhance interaction quality in HRC and foster trust.

The remainder of this paper is structured as follows: Section II describes the methodology; Section III illustrates the experiments; Section IV presents and analyzes the results, which are extensively discussed in Section V, and finally, Section VI concludes the paper.

II. METHODOLOGY

Since human trust is difficult to model, we propose to use PBO that leverages human preferences to guide the optimization process, helping us identify robot behaviors that are most likely to foster trust with human users. We explain the details of our approach in the following sections.

A. Preference-Based Optimization

PBO algorithms iteratively refine a decision vector by relying solely on paired comparisons between different configurations of parameters. Operating as an active preference learning algorithm, each iteration involves a human decision-maker expressing a preference between two options ("this is better than that"), guiding the algorithm toward convergence for the most favored tuning. Among various PBO implementations, the *GLISp* algorithm [13] is selected due to its efficiency in handling the inherent subjectivity of human trust and its capability in solving global optimization problems even when the objective function cannot be directly evaluated. First, a *preference function* $\pi : \mathbb{R}^n \times \mathbb{R}^n \to \{-1, 1\}$ is defined as:

$$\pi(\mathbf{x}_1, \mathbf{x}_2) = \begin{cases} -1 & \text{if } \mathbf{x}_1 \text{ is "better" than } \mathbf{x}_2 \\ 1 & \text{if } \mathbf{x}_2 \text{ is "better" than } \mathbf{x}_1 \end{cases} \tag{1}$$

where $\mathbf{x}_1 \in \mathbb{R}^n$ represents the set of n parameters from the current valuation, while $\mathbf{x}_2 \in \mathbb{R}^n$ denotes the current preferred decision vector. Given the lower and upper bounds $\mathbf{x}_{\min}, \mathbf{x}_{\max} \in \mathbb{R}^n$ for the parameters of the decision vector, the objective is to solve the following preference-based optimization problem:

$$\begin{aligned} \text{find } \mathbf{x}^* \text{ such that } \pi(\mathbf{x}^*, \mathbf{x}) \leq 0, \quad \forall \mathbf{x} \in X, \\ \mathbf{x}_{\min} \leq \mathbf{x} \leq \mathbf{x}_{\max}, \end{aligned} \tag{2}$$

where $\mathbf{x}^* \in \mathbb{R}^n$ is the optimal decision vector. Since the elements x^i (referred to as decision variables) of $\mathbf{x}$ may have different upper and lower bounds, the variables are scaled to range in $[-1, 1]$, as explained in [13].

For each pair $\mathbf{x}_i, \mathbf{x}_j \in X$ with $\mathbf{x_i} \neq \mathbf{x_j}$, the human is asked to express a preference. These preferences are saved in the preference vector:

$$\begin{aligned} \mathbf{p} = [p_1, \ldots, p_M]^T \in \{-1, 1\}^M \\ \text{where } p_h = \pi(x_{i(h)}, x_{j(h)}), \end{aligned} \tag{3}$$

with M indicating the number of expressed preferences, $1 \leq M \leq \binom{N}{2}$, $h \in \{1, \ldots, M\}$ and N is the number of samples.

The algorithm iteratively guides the decision-maker to the global optimum by defining a surrogate of the latent objective function using past decision vectors and pairwise preferences.

In this study, the *surrogate function* $\hat{f} : \mathbb{R}^n \to \mathbb{R}$ was constructed using the *radial-basis function* (RBF):

$$\hat{f}(\mathbf{x}) = \sum_{i=1}^{N} \beta_i \Phi(\varepsilon d(\mathbf{x}, \mathbf{x}_i)), \tag{4}$$

where $d : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ is the Euclidean distance, $\varepsilon > 0$ is the shape parameter of RBF, $\Phi : \mathbb{R} \to \mathbb{R}$ is the RBF $\frac{1}{1+(\varepsilon d)^2}$, and β_i are the unknown weights to be computed based on the available user's preferences. These coefficients are determined by solving the Quadratic Programming problem in [13], with a tolerance parameter $\sigma > 0$ used for the constraints when fitting the surrogate. The performance of GLISp can be enhanced through a *re-calibration of the parameter* ε. This process involves conducting leave-one-out cross-validation on the existing samples to determine the scaling factor $\varepsilon \leftarrow \theta\varepsilon$. With reference to [13], the variable θ represented the vector of values $\{10^{-1+\frac{1}{5}(l-1)}\}_{l=1}^{10}$ used in the re-calibration phase which is performed at the steps indexed

by $I_{sc} = \{N_{init}, N_{init} + \lfloor \frac{N_{max}-N_{init}}{4} \rfloor, N_{init} + \lfloor \frac{N_{max}-N_{init}}{2} \rfloor, N_{init} + \lfloor \frac{3(N_{max}-N_{init})}{4} \rfloor\}$. The number of initial samples N_{init} is selected as $\frac{N_{max}}{3}$, with N_{max} representing the maximum number of allowed evaluations performed by the decision-maker. The initial samples to initialize the algorithm were generated using *Latin Hypercube Sampling* [22], which systematically produces a representative sample set from a multidimensional probability distribution by dividing the parameter space into bins and selecting one point randomly from each bin, ensuring efficient and non-redundant sampling across the distribution.

Relying solely on the surrogate function is insufficient to ensure the discovery of the global minimum. Further exploration of the space is necessary to avoid missing potential optimal solutions. In the GLISp algorithm, an acquisition function is defined to balance exploration (i.e, exploring feasible areas of the domain) and exploitation (i.e., exploiting surrogate to optimize the parameters) when gathering a new sample. The *exploration function* is defined as the *Inverse Distance Weighting* function $z : \mathbb{R}^n \to \mathbb{R}$, modified as in [14], to prioritize increased exploration during the initial iterations and gradually diminish its impact in favor of exploitation as the number N of iterations grows:

$$z_N(\mathbf{x}) = \begin{cases} 0, & \text{if } \mathbf{x} \in \{\mathbf{x}_1, \dots, \mathbf{x}_N\} \\ \left(1 - \frac{N_{max}}{N_{min}}\right) \tan^{-1}\left(\frac{\sum_{i=1}^{N} w_i(\mathbf{x}_N^*)}{\sum_{i=1}^{N} w_i(\mathbf{x})}\right) + & \\ \quad + \frac{N}{N_{max}} \tan^{-1}\left(\frac{1}{\sum_{i=1}^{N} w_i(\mathbf{x})}\right), & \text{otherwise} \end{cases} \tag{5}$$

where $w_i : \mathbb{R}^n \to \mathbb{R}$ is $\frac{1}{d^2(\mathbf{x},\mathbf{x}_i)}$ and $\mathbf{x}_N^*$ represents the preferred decision vector discovered up to iteration N. Knowing both $\hat{f}$ and z_N functions, the *acquisition function* is defined as:

$$a(x) = \frac{\hat{f}(x)}{\Delta\hat{F}} - \delta z_N(x), \tag{6}$$

where $\delta \geq 0$ is the exploration parameter and $\Delta\hat{F} = \max_i\{\hat{f}(x_i)\} - \min_i\{\hat{f}(x_i)\}$ is the range of the surrogate function on the samples. Therefore, the next sample $\mathbf{x}_{N+1}$ is provided by solving the global optimization problem:

$$\mathbf{x}_{N+1} = \underset{\mathbf{x} \in X,\ \mathbf{x}_{min} \leq \mathbf{x} \leq \mathbf{x}_{max}}{\arg\min} \quad a(\mathbf{x}), \tag{7}$$

which the solution is determined using the solver *Particle Swarm Optimization* [23].

B. Robot's Trajectory Adaptation based on Human Preferred Interaction Parameters

To guide the optimization process, the user is invited to express their preferences regarding the trustworthiness of the robot's performance while moving from point **A** to point **B**. In this work, the robot's performance is fine-tuned by acting on three well-recognized variables affecting the user's trust during the interaction: (1) the robot velocity profile, (2) human-robot separation distance, and (3) vertical proximity to the user's head. These variables were formulated as suitable inputs for the PBO algorithm and trajectory adaptation model:

1) the *total execution time* to complete the trajectory (τ_{des}). τ_{des} is more effective than considering each component of the velocity, allowing us to tune the robot's velocity profile according to human preferences operating on a single parameter;
2) the *separation distance* between the robot's end-effector and the human body (d_{des});
3) the *height* of the target pose (h_{des}), which directly determines the vertical proximity of the robot's path to the user's head during the trajectory.

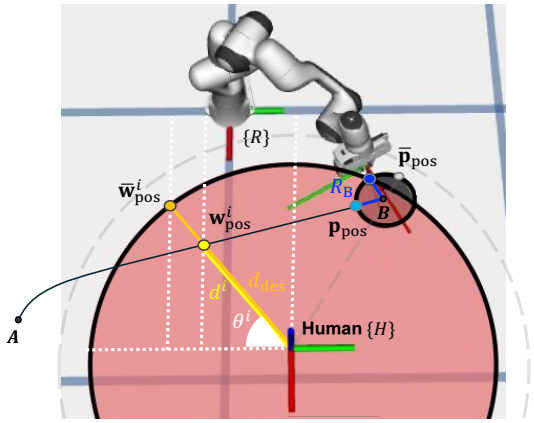

Fig. 1: Robot's trajectory to reach the target and maintain a human-robot separation distance d_{des}. The default and adjusted positions of the i-th waypoint along the trajectory are shown as yellow and orange points, respectively, with the modified target approach point in blue. The angle θ represents the deviation between the radial vector through $\mathbf{w}^i$ and the y-axis. A circle around the target point **B** marks potential pouring locations for the robot.

These *interaction parameters*, constituting the decision vector $\mathbf{x} = [\tau_{des}, d_{des}, h_{des}] \in \mathbb{R}^3$, serve as the decision variables within the PBO algorithm outlined in Section II-A. Consequently, in each iteration of the task, the algorithm provides a new set of parameters, and the robot's trajectory is modified accordingly. The value of τ_{des} generated by the PBO algorithm determines the velocity profile of the robot's end-effector during the performed trajectory. On the other hand, the target approaching location and the path of the robot's end-effector to reach that location are calculated at each iteration based on the values of d_{des} and h_{des}.

More specifically, a circumference, centered at the human's reference frame $\{H\}$, is defined with a radius equal to the distance parameter d_{des}. In each task cycle, the trajectory adaptation strategy prevents the robot from entering the circle and ensures the specified separation distance d_{des} is maintained. During the trajectory execution, the robot transitions to a target pose $\mathbf{B} = [B_x, B_y, B_z]^T \in \mathbb{R}^3$ by following a straight-line path on the horizontal plane, defined by a set of waypoints $\mathbf{w}^i \in \mathbb{R}^6$ with $i \in \{1, \dots, W\}$. Instead of reaching position **B** directly, we assume the robot's end-effector moves to a target approaching pose $\mathbf{p} \in \mathbb{R}^6$ at a distance R_B from **B**, as illustrated in Fig. 1. Note that the z-position coordinate p_z of the target approaching pose is set to the target height h_{des} provided by the PBO algorithm. Along this path, if any waypoints fall inside the circle, their positions are adapted to lie on the circle's edge.

Given the position of the human's reference frame $\{H\}$, denoted as $\mathbf{H}_{pos} = [H_x, H_y, H_z]^T \in \mathbb{R}^3$, the i-th waypoint

$\mathbf{w}^i_{\text{pos}} = [w^i_x, w^i_y, w^i_z] \in \mathbb{R}^3$ is moved along the direction of the vector connecting $\{H\}$ and $\mathbf{w}^i_{\text{pos}}$ on the horizontal plane in order to position it on the circumference of the circle. The height of the i-th waypoint w^i_z is instead modified to gradually reach the desired target height h_{des}. This is implemented by designing a straight line on the vertical yz-plane connecting the first waypoint $[w^1_y, w^1_z]$ with the desired target approaching point $[p_y, h_{\text{des}}]$:

$$\frac{\bar{w}^i_z - w^1_z}{h_{\text{des}} - w^1_z} = \frac{\bar{w}^i_y - w^1_y}{p_y - w^1_y} \tag{8}$$

and sampling it for the modified y-coordinate $\bar{w}^i_y$. In conclusion, the modified position of the i-th waypoint $\bar{\mathbf{w}}^i_{\text{pos}}$ is computed as:

$$\begin{aligned} \bar{w}^i_x &= H_x - d_{\text{des}} \sin\theta^i = H_x - d_{\text{des}} \frac{(H_x - w^i_x)}{d^i}, \\ \bar{w}^i_y &= H_y - d_{\text{des}} \cos\theta^i = H_y - d_{\text{des}} \frac{(H_y - w^i_y)}{d^i}, \\ \bar{w}^i_z &= \frac{(h_{\text{des}} - w^1_z)}{p_y - w^1_y}(\bar{w}^i_y - w^1_y) + w^1_z, \end{aligned} \tag{9}$$

where θ^i is the angle between the radial vector passing through $\mathbf{w}^i$ and the y-axis and d^i is the default distance between $\{H\}$ and $\mathbf{w}^i_{\text{pos}}$.

If the target approaching pose $\mathbf{p}$ of the robot also falls within the protective circle around the human (i.e., the distance d_{HB} between $\{H\}$ and $\mathbf{B}$ is smaller than d_{des}), a new position and orientation $\bar{\mathbf{p}}$ is determined. To achieve this, another circumference is delineated, centered in $\mathbf{B}$ with a fixed radius $R_B \in \mathbb{R}$, indicating the space of potential target approaching locations for the robot. The new target approaching position $\bar{\mathbf{p}}_{\text{pos}} = [\bar{p}_x, \bar{p}_y, \bar{p}_z]$ is computed by solving a system of equations representing the circumferences and determining the intersection points, i.e.,

$$[\bar{p}_x, \bar{p}_y, \bar{p}_z] = \begin{cases} (p_x - H_x)^2 + (p_y - H_y)^2 = d^2_{\text{des}} \\ (p_x - B_x)^2 + (p_y - B_y)^2 = R^2_B \\ \bar{p}_z = h_{\text{des}} \end{cases} \tag{10}$$

Considering the layout of the experimental workspace, $\bar{\mathbf{p}}_{\text{pos}}$ is defined as the point with the lowest x-coordinate, thereby closer to the robot and farther from the person (indicated by a blue point in Fig. 1). Consequently, depending on the value of d_{des}, the intersection point on the circumference surrounding $\mathbf{B}$ changes in each cycle of the task. In case of no intersection (i.e., $d_{HB} < d_{\text{des}}$, as for the dashed circumference in Fig. 1), the robot reaches the farthest possible location from the person. Specifically, the modified target approaching point, depicted by a gray point in Fig. 1, is computed as follows:

$$\begin{aligned} \bar{p}_x &= H_x - (d_{\text{des}} + R_B) \frac{(H_x - B_x)}{d_{HB}}, \\ \bar{p}_y &= H_y - (d_{\text{des}} + R_B) \frac{(H_y - B_y)}{d_{HB}}, \\ \bar{p}_z &= h_{\text{des}}. \end{aligned} \tag{11}$$

Once the target approaching position is defined and reached, the target orientation is determined to align the robot's end-effector perpendicular to the radial vector directed towards $\mathbf{B}$. To achieve this, a rotation on the xy-plane is executed.

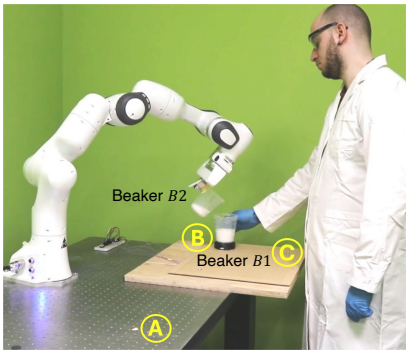

Fig. 2: Experimental setup in chemical industrial scenario: a subject moves beaker *B1* from **C** to **B**, while a robot grasps beaker *B2* at **A** and moves to **B** for pouring a chemical into *B1* held by the subject.

III. EXPERIMENTS

This section presents the experimental procedure to evaluate the effectiveness of the proposed framework in improving human trust toward the robot by using the PBO algorithm. The following research questions (RQ) were investigated:

RQ1. *How effective is the PBO in aligning the robot's interaction parameters with human preferences?*

RQ2. *How does the optimal parameter combination obtained by PBO (denoted as PBO-based parameters) differ from a set of parameters obtained through individual selection (IS-based parameters) of each of them?*

RQ3. *In cases of disparity between PBO-based and IS-based parameters, which approach fosters increased human trust towards the robot?*

A. Experimental Setup and Task Description

The study involved 14 healthy participants (8 males and 6 females, age: 29.21 ± 4.90 years) recruited from the students and research personnel of the Istituto Italiano di Tecnologia[1]. Participants were categorized into two groups based on their expertise in HRC: robotics experts, who possessed prior experience programming a manipulator, and novices, who lacked such experience. The experimental setup, illustrated in Fig. 2, involved a collaborative robot (Franka Emika Panda, controlled at 1kHz and equipped with a two-fingered parallel gripper) and a human subject. The architecture and communication protocols were implemented through the Robot Operating System (ROS) platform.

This study considered a cyclic HRC task, simulating a chemical industrial environment. The participant placed beaker $B1$ at location (**C**), while the robot grasped beaker $B2$ situated at (**A**), which contained a chemical[2]. Upon a sound signal, the participant moved the beaker $B1$ from **C** to a designated point (**B**) and held it there. Meanwhile, the robot moved beaker $B2$ toward **B** and poured the chemical into beaker $B1$ held by the subject. In the pouring phase, the robot's end-effector underwent two sequential rotations. The

[1]Experiments were carried out following the Helsinki Declaration, and the protocol was approved by the ethics committee ASL3 Genovese (Protocol IIT_HRII_ERGOLEAN 156/2020).

[2]The participants were told that the chemical was harmful to make the task more realistic but coarse salt was employed to ensure participant safety.

TABLE I: GLISp Experimental Parameters.

Parameter	Value	Description
$\tau_{\min}, d_{\min}, h_{\min}$	$\{5, 0.520, 0.260\}$	Parameters' lower bounds
$\tau_{\max}, d_{\max}, h_{\max}$	$\{10, 0.650, 0.400\}$	Parameters' upper bounds
N_{max}	15	Max number of valuations
N_{init}	$N_{\max}/3 = 5$	Initial samples
δ	1.0	Exploration parameter
RBF Function	*Inverse quadratic*	To fit the surrogate function
Initial ε	1.0	Shape parameter of RBF
I_{sc}	$\{5, 7, 9, 12\}$	Steps when recalibrating ε
θ	$\{10^{-1+\frac{1}{5}(l-1)}\}_{l=1}^{10}$	Values to recalibrate ε
σ	$1/N_{\max} = 0.06$	Tolerance parameter

former occurred within the *xy*-plane to orient $B2$ for pouring along the radial vector directed towards $B1$ in the determined pouring position $\mathbf{p}_{\text{pos}}$. The latter was a pre-defined rotation in the *xz*-plane to actually pour the chemical. Finally, the robot returned to **A** to release beaker B2, and the subject moved back beaker $B1$ to **C**.

B. Experimental Parameters

The *GLISp* algorithm was fine-tuned and configured with the parameters illustrated in Table I. In this study, the *total execution time* $\tau_{\text{des}}[s]$ represented the time taken for the robot to move $B2$ to the pre-pouring location $\mathbf{p}_{\text{pos}}$ (equal to the time taken to come back in **A**)[3]. The *separation distance* $d_{\text{des}}[m]$ was computed as the distance between the human pelvis reference frame and the robot's end-effector. Finally, the *height* $h_{\text{des}}[m]$ corresponded to the height at which the pouring was executed. To interpolate the adapting waypoints described in Sec. II-B, we employed *Quintic B-splines* curves that allow smooth trajectories featuring continuous derivatives up to the jerk. For further details about B-splines, please refer to [12]. The trajectory planner sent then appropriate references to a low-level *Cartesian impedance controller*.

C. Experimental Protocol

The experimental protocol included two phases.

1) Preferred Parameters Selection Experiment: The first phase was conducted to acquire users' preferences in terms of interaction parameters for the HRC task and thus investigate RQ1 and RQ2. Two methods were considered in this phase:

(1a) ***PBO****:* Participants engaged in the task 16 times, with feedback collected after each task cycle, yielding $N_{max} = 15$ evaluations. Specifically, at the beginning of the experiment, 2 cycles were executed with different sets of parameters and the subjects expressed their preference (*feedback*) between them. Then, at the end of each cycle, they expressed a preference between the set of parameters just experienced and the preferred one up to that cycle. When necessary, a subject could replicate a cycle with the preferred set. To express their preference, participants used a hand gesture. After each human feedback, the optimization algorithm was updated until the last cycle, returning the final PBO-based combination of the 3 interaction parameters.

(1b) ***IS****:* This approach involved 9 repetitions of the task divided into 3 sets of 3 task iterations. For each set, 2 interaction parameters were constant (median value of the respective range) while the third one varied (lower bound of the range at the first iteration, median value at the second iteration, or higher bound at the third iteration). For example, in the first set, the distance and pouring height were set to their median value while the execution time varied. At the end of each set, participants selected one of the 3 values (low, medium, high) of the varying parameter. After the 3 sets, the selected parameter values formed the IS-based set.

All the participants completed both methods (i.e., *1a* and *1b*), with the order randomized and a break in between to avoid fatigue. At the end of phase *1a*, subjects were instructed to rate the importance assigned to each interaction parameter in fostering trust in the robotic assistant (i.e., respond to question *Q1* of the questionnaire[4]) using a 5-point Likert scale. The responses of each subject s were normalized in $[0, 1]$ and stored in the variables $\mathrm{I}_{\tau}^{s}, \mathrm{I}_{d}^{s}$, and I_{h}^{s}.

2) Preferred Parameters Evaluation Experiment: The second phase was conducted to investigate RQ3, namely determining which set of parameters (PBO-based versus IS-based) fosters increased human trust toward the robot. Participants were asked to perform 3 consecutive cycles of the HRC task in 2 conditions: *(2a)* with the PBO-based set of parameters (output of phase *1a*) and *(2b)* with the IS-based one (output of phase *1b*). The order of the conditions was randomized and a break in between was introduced. After each condition (i.e. *2a* and *2b*), participants were required to complete the second section of the questionnaire[4] (see questions *Q2* to *Q4* in Table II). Note that *Q3* encompassed 5 items from *Schaefer's Trust Questionnaire* [24]), while *Q4* entailed 7 items from *Charalambous's Trust Questionnaire* [25]. Responses were normalised to obtain values $\in [0, 1]$.

D. Experimental Analysis

In the first experimental phase (phase *1a*), we evaluate the performance of PBO by analyzing the **evolution** and **convergence** of the model across iterations, addressing RQ1. Parameter convergence was defined as reaching a boundary within *one-third* of the corresponding interval range from the final selected value and maintaining stability within that bound thereafter. Furthermore, we normalized each PBO-obtained parameter within its range and assessed the standard deviation across participants to explore the necessity of parameter personalization for fostering trust.

To address RQ2, we examined the **disparity** between results from the PBO algorithm and individual selection of parameters (PBO-based versus IS-based). For each participant, we calculated and normalized the differences between PBO-based and IS-based parameter values across their respective ranges. However, it is crucial to acknowledge that each parameter might hold different significance in shaping participants' trust levels and comfort. For instance, a participant may prioritize the velocity profile of the robot over

[3]To facilitate the users in appreciating the differences between different values of total execution time, the duration of the transition phases (e.g., alignment with target pose, pouring, etc.) was tuned as a function of τ_{des}.

[4]The overall questionnaire can be found at: https://www.campagna-robotics.com/_files/ugd/1d4f1d_1d1d1e4594e24225a9eaed7692caf334.pdf

other parameters, thereby choosing the robot's trajectories accordingly. To account for this variability, we weighted each parameter difference based on the importance assigned by the specific participant to that parameter (i.e., I^s_τ, I^s_d, and I^s_h, computed from question *Q1*) and computed the mean value. Thus, the overall disparity Δ^s between PBO-based and IS-based parameters for subject s was computed as follows:

$$\Delta^s = \frac{1}{I^s_\tau + I^s_d + I^s_h}\Big[I^s_\tau \frac{|\tau^s_{\text{PBO}} - \tau^s_{\text{IS}}|}{\tau_{\max} - \tau_{\min}} + I^s_d \frac{|d^s_{\text{PBO}} - d^s_{\text{IS}}|}{d_{\max} - d_{\min}} + I^s_h \frac{|h^s_{\text{PBO}} - h^s_{\text{IS}}|}{h_{\max} - h_{\min}}\Big], \tag{12}$$

where $\tau^s_{\text{PBO}}, d^s_{\text{PBO}}, h^s_{\text{PBO}}$ denote the PBO-based optimal parameters for subject s and $\tau^s_{\text{IS}}, d^s_{\text{IS}}, h^s_{\text{IS}}$ the IS-based ones.

Finally, a statistical analysis was conducted to compare the perceived trust reported in the questionnaire across testing conditions *2a* and *2b*, thereby addressing RQ3. The results of *Q2*, *Q3*, and *Q4* were initially checked for normality using the *Shapiro-Wilk test*. Once normality was confirmed, *paired t-tests* were performed. Otherwise, we used the non-parametric *Wilcoxon signed-rank test* to evaluate any significant differences.

IV. EXPERIMENTAL RESULTS

1) Preferred Parameter Selection Experiment: Fig. 3 illustrates the evolution and convergence of each parameter during PBO (phase *1a*) for a participant. Blue circles represent the values suggested by the optimization algorithm at each iteration, while red crosses denote user-expressed preferences between the current iteration's value and the preferred value in previous iterations. A yellow band marks the convergence bound for each parameter, and a red asterisk denotes the optimal PBO-based parameter value.

Fig. 4 summarizes the results for all participants. The PBO algorithm effectively explored the parameters' search space, as indicated by user preferences (red dots) spanning each parameter's range across iterations, converging to the final preferred value (marked by a red asterisk). Across participants, all parameters reached convergence with the specified preferences after an average of 12.9 ± 1.8 iterations from the beginning of the interaction. Interestingly, participants showed diverse preferences, with notable variations observed across all 3 parameters. Indeed, normalizing PBO-based parameters within their ranges revealed average standard deviations of

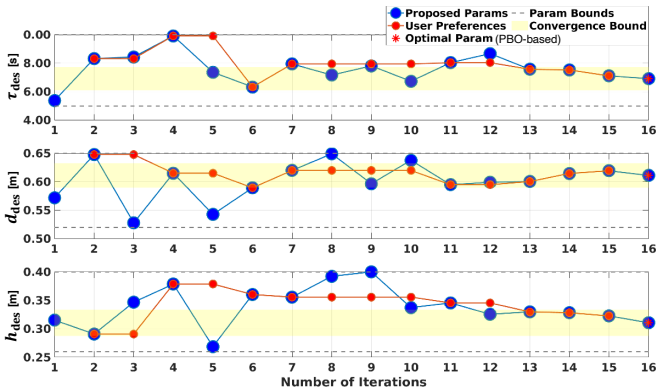

Fig. 3: PBO evolution until convergence for subject 11.

0.413 for execution time, 0.345 for separation distance, and 0.397 for pouring height across participants.

To compare PBO-based and IS-based parameters, the latter are depicted in Fig. 4 via gray bounds, whose varying intensities represent the importance assigned to each parameter by the specific participant. We can observe that the red asterisk (PBO-based parameters) fell within the gray box (IS-based ones) in more than half of the cases. Considering only the most trust-relevant parameter for each participant, the PBO-based preference matched the IS-based one for 10 out of 14 participants, resulting in an overall alignment of 71.43%. At the bottom of the figure, the overall disparity Δ^s obtained for each participant using Equation (12) is presented. Significant disparities, defined as those exceeding one-third of the maximum potential disparity, are indicated with a star.

2) Preferred Parameter Evaluation Experiment: Table II reports the mean and standard deviation of trust scores derived from the overall questionnaire responses while the robot utilized PBO-based parameters. Upon comparison with outcomes obtained through the use of IS-based parameters, no statistically significant differences were observed. Therefore, we conducted a detailed analysis of individual participant responses and preferences to gain deeper insights. Table III presents a subject-by-subject comparison of trust scores between PBO-based and IS-based parameters. To aid in-

TABLE II: Results of the questionnaire with PBO-based parameters for all subjects. For *Q2*, scoring 0 indicates that the subject strongly disagrees and 1 strongly agrees with the statement, while in *Q3* and *Q4* the higher the score the higher the perceived trust in HRC.

Questions	Mean	Std
Q2: The way the robot moved during this experimental session promoted my trust in the robot.	0.82	0.18
Q3: Schaefer's Trust Questionnaire	0.79	0.22
Q4: Charalambous Trust Scale		
a. Perceived robot motion and pick-up speed	0.78	0.15
b. Perceived cooperation safety	0.82	0.17
c. Perceived robot reliability	0.78	0.23

TABLE III: Subject-by-subject comparison of trust scores for questions in Table II between PBO-based and IS-based parameter adoption, specifying users' expertise in HRC.

Subject	Experience	*Q2*	*Q3*	*Q4*		
				a	*b*	*c*
1	expert					
2	expert					
3	expert					
4	expert					
5	novice					
6	novice					
7	novice					
8	novice					
9	novice					
10	novice					
11	expert					
12	expert					
13	expert					
14	expert					

Higher trust score with PBO
Higher trust score with IS
Equal trust score

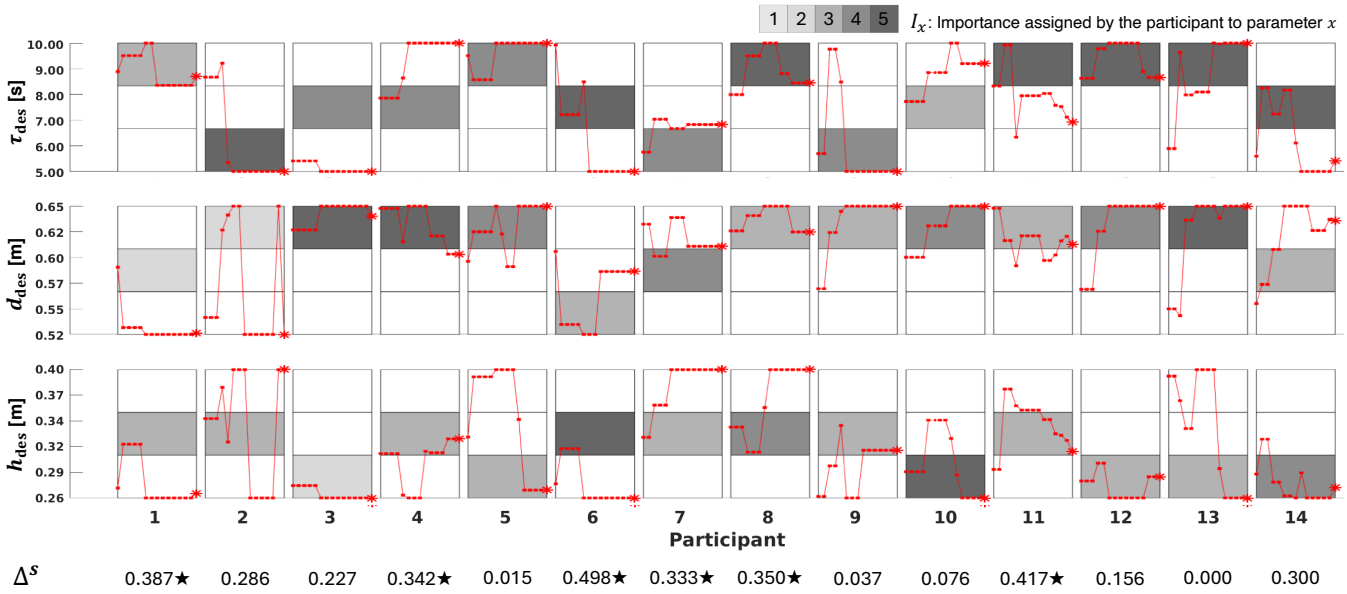

Fig. 4: Optimal parameters' combination obtained by PBO (PBO-based parameters, denoted as red asterisks) for all participants and comparison with individual selection of each parameter (IS-based parameters, represented by gray boxes). Gray intensity reports the importance assigned to the specific parameter by the participant in *Q1* of the questionnaire. Red dots represent the user's preferences expressed throughout the iterations. At the bottom, the overall disparity indexes Δ^s are reported. Significant ones are marked with a star.

terpretation, we highlighted higher trust scores using PBO-based parameters in violet, IS-based preferences in pink, and equal trust scores in yellow. Participant experience in HRC is also reported. Robotics novices generally favored PBO-based parameters, evident from the prevalence of violet cells. Overall, novices' trust in the robot improved across both custom and standard questions, with growth rates observed at 18.35%, 10.23%, and 10.71% for *Q2*, *Q3*, and *Q4*, respectively.

V. DISCUSSION

The analysis of PBO evolution over time demonstrated its effectiveness in reducing robot behavior variability and achieving stable interaction parameters aligning with individual preferences. As shown in Fig. 3, the algorithm initially explores the parameter space extensively, then gradually adopts a more exploitative approach until solutions remain within the convergence region (yellow band). Successful convergence occurred for 13 out of 14 participants, indicating consistent performance. However, for subject 2, the algorithm struggled to find the optimal parameter combination, likely due to the user's difficulty in perceiving changes in specific robot's interaction parameters, as also evidenced by inconsistent questionnaire responses (see Table III). The final PBO results (i.e., PBO-based parameters) varied significantly among participants (see Fig. 4), highlighting the subjective nature of trust and the nuanced factors influencing it. This underscores the need to tailor the HRC to individual preferences. Notably, participants differed in which parameter they found most influential on their trust and showed diverse preference evolution over time. Overall, the trajectory's execution time had the most significant impact on trust, as seen in Fig. 4. However, all interaction parameters were rated as most relevant by at least one or more participants. Additionally, we noted that participants who reported the unbalanced importance of the interaction parameters focused on directing the PBO algorithm to optimize a specific subset of parameters according to their priorities, often neglecting the remaining parameters. This tendency resulted in preferences that spanned a wide spectrum of the range, as observed for subject 2.

The comparative analysis between PBO-based and IS-based parameters revealed that when focusing solely on the "most influential to trust" parameter for each participant, the PBO algorithm closely matched the IS-based preferences in most of cases (71.43%). This finding suggests that participants possess a full awareness of how this highest-importance parameter value impacts robot movements and, consequently, their own perceived trust. On the other hand, when considering all interaction parameters, differences emerged in users' preferences between the two approaches. Indeed, a notable overall disparity, denoted as Δ^s, was registered for almost half of the subjects. This inconsistency in preferences expressed by participants using the PBO and IS implies that participants may not be able to directly identify their preferred values for all the parameters. Thus, using an indirect and implicit approach, such as through PBO, could potentially be beneficial in personalizing interaction parameters.

Subjective questionnaire responses indicated that robot movements utilizing PBO-based parameters significantly enhanced trust in the robotic partner. Statistical analysis across all participants showed no significant difference compared to IS-based parameters. However, closer inspection revealed PBO was more effective for novice users in identifying optimal parameter combinations to enhance trust. Specifically, 4 out of 6 novices perceived greater trust when the robot used PBO-based parameters versus IS-based ones. This indicates novices may find it challenging to express preferences for

interaction parameters independently due to their limited expertise in HRC. Indeed, novices may struggle to grasp how parameter combinations influence robot movements and subsequently affect trust, and they are likely more sensitive to changes in robot behavior. Conversely, the trust level of robotic experts did not seem to be impacted by interaction parameter changes, with reported incoherent trust scores across items (alternance of violet and pink boxes in Table III for the same subject) or even a preference for the IS-based parameters in 2 out of 8 participants (see subjects 3 and 13 in Table III). Experienced individuals in robot programming likely have clearer expectations, thus finding IS-based selections more effective.

VI. CONCLUSION

The study introduced a novel framework aimed at fostering an adequate level of user trust toward the robot by optimizing robot interaction parameters and adapting trajectories accordingly. Given the inherently subjective nature of trust, PBO emerged as a suitable optimization algorithm, as it considers parameters collectively and relies on qualitative feedback provided by humans to suggest a refined set of parameters for subsequent iterations. The algorithm and resulting trajectory adaptation were effective in enhancing users' confidence in HRC, particularly for robotics novices. In contrast, individuals with prior experience in robotics showed similar trust perceptions as with direct parameter-by-parameter selection. The framework was designed to be versatile and applicable to various HRC tasks involving close human-robot proximity, with only the need for adaptations to specific task requirements (e.g., trajectory waypoints).

The research was limited by the fixed-base robot's restricted workspace, which constrained the range of parameter changes. To address this, a new setup with a wider range of motion will be developed, enabling broader trajectory adaptations to elicit diverse trust responses. Future work will analyze individual sensitivity to parameter variations to tailor ranges accordingly. Additionally, more advanced interaction parameters and alternative feedback mechanisms will be explored to enhance robot-human understanding and insights into human trust variations. Moreover, this study was designed as a preliminary investigation, and future research will involve a larger pool of participants. Finally, occasional contradictory findings raise doubts about the reliability of questionnaires and encourage us to investigate trust indicators derived from human-body language and emotional cues presented in the literature [19].

REFERENCES

[1] S. Panagou, W. P. Neumann, and F. Fruggiero, "A scoping review of human robot interaction research towards industry 5.0 human-centric workplaces," *International Journal of Production Research*, vol. 62, pp. 974–990, 2024.

[2] B. M. Muir and N. Moray, "Trust in automation. part ii. experimental studies of trust and human intervention in a process control simulation," *Ergonomics*, vol. 39, pp. 429–460, 1996.

[3] J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," *Human factors*, vol. 46, pp. 50–80, 2004.

[4] K. Hald, M. Rehm, and T. B. Moeslund, "Human-robot trust assessment using top-down visual tracking after robot task execution mistakes," in *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*. IEEE, 2021, pp. 892–898.

[5] L. Onnasch and C. L. Hildebrandt, "Impact of anthropomorphic robot design on trust and attention in industrial human-robot interaction," *ACM Transactions on Human-Robot Interaction (THRI)*, vol. 11, pp. 1–24, 2021.

[6] R. van den Brule, R. Dotsch, G. Bijlstra, D. H. Wigboldus, and P. Haselager, "Do robot performance and behavioral style affect human trust? a multi-method approach," *International journal of social robotics*, vol. 6, pp. 519–531, 2014.

[7] C. Li, X. Zhang, D. Chrysostomou, and H. Yang, "Tod4ir: A humanised task-oriented dialogue system for industrial robots," *IEEE Access*, vol. 10, pp. 91 631–91 649, 2022.

[8] B. F. Malle and D. Ullman, "A multidimensional conception and measure of human-robot trust," in *Trust in human-robot interaction*. Elsevier, 2021, pp. 3–25.

[9] M. Rueben, S. A. Elprama, D. Chrysostomou, and A. Jacobs, "Introduction to (re) using questionnaires in human-robot interaction research," *Human-Robot Interaction: Evaluation Methods and Their Standardization*, pp. 125–144, 2020.

[10] M. Lagomarsino, M. Lorenzini, M. D. Constable, E. De Momi, C. Becchio, and A. Ajoudani, "Maximising coefficiency of human-robot handovers through reinforcement learning," *IEEE Robotics and Automation Letters*, pp. 1–8, 2023.

[11] C. Messeri, G. Masotti, A. M. Zanchettin, and P. Rocco, "Human-robot collaboration: Optimizing stress and productivity based on game theory," *IEEE Robotics and Automation Letters*, pp. 8061–8068, 2021.

[12] M. Lagomarsino, M. Lorenzini, E. De Momi, and A. Ajoudani, "PROMIND: Proximity and reactivity optimisation of robot motion to tune safety limits, human stress, and productivity in industrial settings," *IEEE Transactions on Robotics (T-RO)*, 2024, (Under revision).

[13] A. Bemporad and D. Piga, "Global optimization based on active preference learning with radial basis functions," *Machine Learning*, vol. 110, pp. 417–448, 2021.

[14] M. Zhu, D. Piga, and A. Bemporad, "C-glisp: Preference-based global optimization under unknown constraints with applications to controller calibration," *IEEE Transactions on Control Systems Technology*, vol. 30, pp. 2176–2187, 2021.

[15] L. Roveda, B. Maggioni, E. Marescotti, A. A. Shahid, A. M. Zanchettin, A. Bemporad, and D. Piga, "Pairwise preferences-based optimization of a path-based velocity planner in robotic sealing tasks," *IEEE Robotics and Automation Letters*, vol. 6, pp. 6632–6639, 2021.

[16] M. Tucker, E. Novoseller, C. Kann, Y. Sui, Y. Yue, J. W. Burdick, and A. D. Ames, "Preference-based learning for exoskeleton gait optimization," in *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2020, pp. 2351–2357.

[17] E. Bıyık, N. Huynh, M. J. Kochenderfer, and D. Sadigh, "Active preference-based gaussian process regression for reward learning and optimization," *The International Journal of Robotics Research*, vol. 43, pp. 665–684, 2024.

[18] M. Lorenzini, M. Lagomarsino, L. Fortini, S. Gholami, and A. Ajoudani, "Ergonomic human-robot collaboration in industry: A review," *Frontiers in Robotics and AI*, p. 262, 2023.

[19] M. Lagomarsino, M. Lorenzini, P. Balatti, E. De Momi, and A. Ajoudani, "Pick the right co-worker: Online assessment of cognitive ergonomics in human-robot collaborative assembly," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 15, pp. 1928–1937, 2023.

[20] G. Campagna and M. Rehm, "Analysis of proximity and risk for trust evaluation in human-robot collaboration," in *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2023, pp. 2191–2196.

[21] K. Hald, M. Rehm, and T. B. Moeslund, "Proposing human-robot trust assessment through tracking physical apprehension signals in close-proximity human-robot collaboration," in *2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2019, pp. 1–6.

[22] M. D. McKay, R. J. Beckman, and W. J. Conover, "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code," *Technometrics*, pp. 55–61, 2000.

[23] A. I. F. Vaz and L. N. Vicente, "Pswarm: a hybrid solver for linearly constrained global derivative-free optimization," *Optimization Methods & Software*, vol. 24, pp. 669–685, 2009.

[24] K. E. Schaefer, "Measuring trust in human robot interactions: Development of the "trust perception scale-hri"," in *Robust intelligence and trust in autonomous systems*. Springer, 2016, pp. 191–218.

[25] G. Charalambous, S. Fletcher, and P. Webb, "The development of a scale to evaluate trust in industrial human-robot collaboration," *International Journal of Social Robotics*, vol. 8, pp. 193–209, 2016.