User-Aware Collaborative Learning in Human-Robot Interactions

Balint Gucsi[1], Nguyen Tan Viet Tuyen[1], Bing Chu[1], Danesh Tarapore[1] and Long Tran-Thanh[2]

Abstract— **Our work investigates how social robots can efficiently collaborate with human users in a user-aware manner, minimising the generated frustration in human colleagues, thus enhancing their experience. As part of this, we develop a user-aware framework for human-robot collaborative learning. We model users' frustration during human-robot interactions based on recent interactions inspired by Psychological principles and develop different frustration-aware interactive preference learning and decision-making models using multi-armed bandit and knapsack methods. Evaluating our approach, 1) we conducted simulated experiments on realistic human-behaviour datasets and 2) a user-study in which participants worked with a TIAGo Steel humanoid robot on a collaboration task using frustration-aware and non frustration-aware (Upper Confidence Bounds and Instruction-based) models. We demonstrate that when collaborating with the frustration-aware robot, users completed the collaboration task 9.04% faster and using 20.54% less number of verbal interactions, with user questionnaire responses reporting less frustration experienced compared to the baseline approaches. Additionally, we create a multimodal dataset containing over 6 hours of human-robot interactions displaying various explicit and implicit user responses.**

I. Introduction

With the development of Artificial Intelligence (AI), robots are deployed in various interaction settings, however, most commonly in isolation from humans, in physically separated areas (e.g. in factories, warehouses). Despite various researches developing efficient systems for human-robot collaboration, the usage of such systems still remains limited in real world settings due to various user concerns [1]. Recent researches and surveys [2], [3], [4] argue that the key to the successful deployment of robotic systems in social settings is the acceptance and trust of human users. The acceptance of human users is commonly influenced by their experience of working with robots – whether they were happy or disappointed by the robot's performance, or whether they felt like the robot actually helped their work or if it only distracted or frustrated them, thus hindering their overall productivity [5]. Apart from the commonly targeted objective metrics such as efficiency or robustness of performing tasks, user experience is heavily influenced by aspects of the interaction such as helpfulness, or predictability [6]. Notably, AI research [7] has pointed out that these factors may have a larger effect on user experience than model accuracy.

In our work, we focus on designing a learning framework for human-robot collaborations which also accounts for the user's frustration and overall experience, in addition to maximising accuracy. Doing so, we address questions such as is it possible to infer how users perceive robotic help or which robotic assistance approaches they find preferable, what they find annoying; and what actions should a robot make to best assist an individual considering their current state. Overall, we seek to answer how to learn to collaborate with users, without unnecessarily exposing them to frustration.

We investigate a scenario where the collaborating parties (i.e. robot and user) are working at the same location, making continuous interactions and feedback possible. As part of this, we develop a frustration-aware learning framework, learning how to interact with individuals by balancing actions and queries throughout a collaboration while learning their preferences without unduly causing them frustration. We assume a collaborative task where a fixed set of actions are performed by the human-robot pair in some sequence (e.g. preparing a dish, assembling furniture or performing household chores) in a leader-follower manner, with the robot assuming the role of an assistant. Throughout the interaction, we model the user's expected frustration level following patterns established in Psychology research, and propose three different user-aware models. The proposed system is evaluated in different interaction scenarios using quantitative experiments on datasets of typical user behaviours and in a qualitative user study of a real Human-Robot collaborative manipulation scenario with the TIAGo humanoid robot [8].

II. Background

Modelling user behaviour and user affective states have a key role in enabling user-aware human-robot collaborations. Involving affective states and their respective social signals in AI models was introduced in Human-Computer Interaction literature [9] and Social signal processing [10], with goals of designing computer systems that users enjoy using. In Human-Robot Interaction (HRI) settings, several works [11], [12], [4], [13] focus on identifying the users' affective states while they participate in HRI scenarios. They primarily rely on user self-assessment to elicit information on the user's experience, affective states (mainly frustration and trust), and the correlation between affective states [4]. Some approaches [4], [14] also attempt to map observed user behaviour to the self-assessment's results.

Other approaches focus on implicit feedback. From the different behaviours exhibited by individuals when experiencing frustration, involuntary bodily functions such as pulse or respiration [15][16], electrocardiogram (ECG) [15], electromyography (EMG) [17] are the most reliable indicators of the affective state. However, due to the complicated sensory requirements of these methods, these are more suitable for laboratory or medical use cases than everyday scenarios.

[1]School of Electronics and Computer Science, University of Southampton, UK {bg1u17,tuyen.nguyen,b.chu,d.s.tarapore}@soton.ac.uk

[2] Department of Computer Science, University of Warwick, UK long.tran-thanh@warwick.ac.uk

This work was supported by UK Research and Innovation [EP/S024298/1]. The authors acknowledge the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton, in the completion of this work.

User affective signals have been successfully inferred using natural language inputs, such as the presence of user frustration (as a binary factor) or future expected frustration based on the user's behaviour when interacting with "smart" tutoring applications [18], [19], [20]; general computer applications [21]; or information retrieval systems [22].

A common theme among inference models in HRI contexts is identifying erroneous events throughout the interaction scenario [23], [24], [25], such as when the robot makes a technical error or performs actions inconsistent with the user's expectations. These models treat the input signals as implicit feedback and do not directly model the affective states experienced by the user. The majority of inference models rely on using the user's facial expressions [20], [22], [18], [25], often as part of multimodal data-streams (e.g. gaze [24], pose [19], gestures [23] or thermal face imaging [26]).

Other approaches focus on inferring the user's impression of the current interaction scenario, interpreting that as a user-provided reward function [27] or as evaluative reinforcement feedback [28]. While doing so, the inference model classifies the user's feedback as positive or negative, or attempts to use the input to generate a ranking, without accounting for the user's underlying affective state. In an attempt to make implicit feedback based decision-making more robust focusing on specific events (e.g. robot error), the approach of combining it with explicit user feedback [29] has been proposed. In contrast, the work of [30] uses facial expressions and speech information of individuals to infer their affective state expressed as an arousal-valence vector, which is then used to determine the robot's interaction behaviour based on its associated intrinsic mood. The model is deployed in a HRI conversation scenario, negotiating a resource splitting problem with an individual, thus the nature of feedback signals differs from the other discussed approaches.

III. Problem Definition

We define our problem as a sequential decision-making problem, with the aim of learning the user's preferred way of executing a multistep collaboration task (and successfully provide assistance with that), while constrained by the user's individual frustration tolerance level. It can be described as $\langle U, T, A, B(\cdot), \mathcal{B}\rangle$, where U is a set of human users, T is set of discrete time slots, A is a set of actions the agent can perform to interact with the user (e.g. assist in certain scenarios, or request additional information), $B(\cdot)$ is the frustration cost function, associating individual actions with a cost and $\mathcal{B} > 0$ is the user's frustration tolerance limit.

Problems of this type may be found in various HRI contexts where the user's preference has to be taken into account to assist efficiently, and this information has to be learnt as part of the interaction as the information is not available in advance (making pre-interaction preference elicitation strategies unsuitable).

IV. Frustration aware learning framework

User Preference Tensor: As part of the learning model, we define a 3-dimensional User Preference Tensor $\mathcal{M}$, representing the likelihood of user $u \in U$ to request assistance with task $a \in A$ at time $t \in T$. Rows of the tensor represent possible actions a and columns represent times t, with probability values $[0, 1]$. We consider the User Preference Tensor the underlying belief model of all users' preferences participating in the interaction scenarios. Regarding individual instances, the user's preferred actions for each time t can be considered to be sampled from $\mathcal{M}$. In practical scenarios, these values are the oracle preferences of real life users or values extracted from user behaviour datasets.

Action Space: When interacting with users, the agent's action space consists of assistive and communication actions. *Assistive actions*: manipulation tasks aimed to assist the user in their current activity (e.g. provide the next necessary ingredient when cooking).
Communication actions: verbal interactions with the user, with the aim of acquiring preference information (i.e. user preference on being assisted with the current task), seeking feedback or exchanging information. Our main communication actions are instruction queries and confirmation queries.

Frustration Cost Model: Our learning framework is designed to support various different frustration cost models. For the purpose of our experiments we propose a frustration model supported by Psychology research, inferring the estimated frustration level from previously observed events (e.g. actions and errors that took place during the interaction).

We consider two main sources of frustration originating from the robot's decisions: interruption[1] and error[2].
Interruption-based Frustration: Interruption-based frustration has been modelled in the past [32], [33] using the so called Bother Cost Model. Since this model has been extensively used in literature and is suitable for modelling users with different characteristics, typical of HRI scenarios, we use it in our work. The interruption-frustration cost is defined as: $FSF_Q = \sum_{q \in Q} c(q)\beta^{t(q)}$ where Q represents past queries, and $c(q)$ the base frustration cost of an interruption (i.e. query) signifying the event's severity (cognitive cost) and duration (time to resolve interruption) and $0 < \beta \leq 1$ representing a discount factor used to simulate the interaction's diminishing impact (i.e. an interruption that happened a longer time ago influences the individual's frustration less than a recent interruption) and $t(q)$ quantifies the amount of time elapsed since the interruption. The total "interruption-based frustration" (IF) level is calculated as:

$$IF = \frac{1 - \alpha_Q^{FSF_Q}}{1 - \alpha_Q} \tag{1}$$

where $\alpha_Q = 1.26 - 0.05w$ with $w \in [1, 10]$ denoting the user's willingness (with 0 being unwilling and 10 being

[1]Interruptions in the form of verbal interactions or questions may disrupt the individual's focus, distract them from their current task, and may demand some cognitive effort on their part to respond appropriately. As such, by the definition of frustration [31], these interruptions can induce frustration in the individual.

[2]Inevitably, as part of the robot's attempts to assist the individual in their task, some attempts will not match the user's expectations. Assistive actions may not suit the current scenario (e.g. the robot offered to vacuum the floor when the individual wants to relax) or the user's preferences (e.g. the robot provided ingredients in the wrong order when cooking), yielding frustration. Naturally, other types of errors can also occur throughout interactions, such as accidental technical errors, communication errors or social errors. However, since these errors do not originate from the agent's decision-making process, rather than the execution of motions, they are out of our scope.

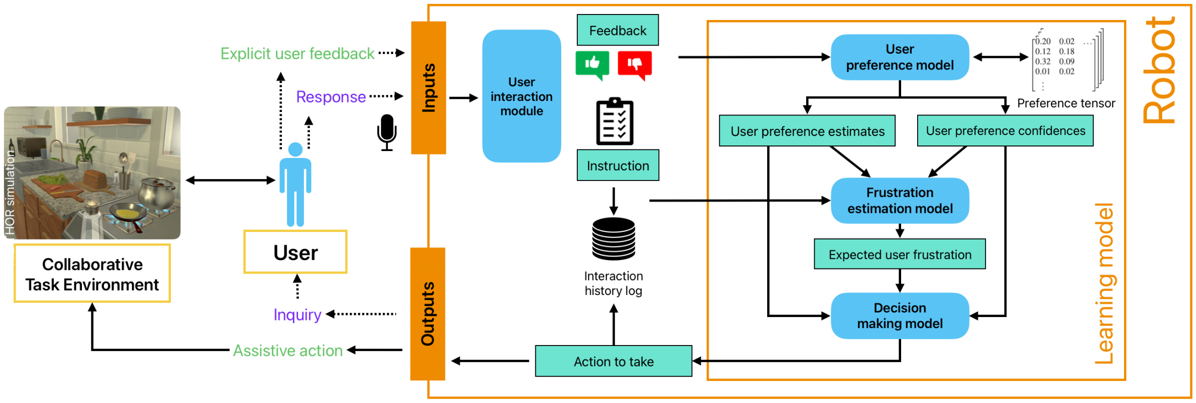

Fig. 1. Overview of the frustration aware user learning framework's components.

most willing). The function is modelled to follow a more exponential or logarithmic pattern for less or more willing users, respectively.

Error-based Frustration: To model error-based frustration, we rely on characteristics of frustration established in the Psychology literature. Specifically, we utilise implications of the Frustration-Aggression Hypothesis [31], [34], Goal-Gradient Principle [35], [36] and the role of Goal Attractiveness [9], [37] regarding user frustration.

Similarly to the model defined by [32], we introduce the "error frustration so far" (FSF_E) function defined as: $FSF_E = \sum_{e \in E} g(e)c(e)\beta^{t(e)}$ where E represents past errors, $0 < \beta \leq 1$ the same discount factor illustrating the error-based frustration's diminishing effect, $c(e)$ representing the frustration's base cost characterised by the error's severity (i.e. difficulty of correcting the error, the permanence of its result) and $g(e)$ defining the goal-gradient scaling factor.

$$g(e) = \gamma + \frac{1-\gamma}{|T|-1} d(e) \tag{2}$$

where $0 < \gamma \leq 1$ is a scaling parameter of the Goal-Gradient Principle ranging between constant and linear (characterised by the goal's importance) and $d(e)$ quantifies the distance until the task's expected completion at the time of the error. Following the Frustration-Aggression Hypothesis and the role of Goal Attractiveness, we calculate the total error-based frustration

$$EF = \frac{1 - \alpha_E^{FSF_E}}{1 - \alpha_E} \tag{3}$$

with $\alpha_E = 1.2 - 0.02p$ with $1 < p \leq 10$ denoting the user's patience or tolerance towards the robot, where a more patient individual experiences less frustration as a result of a logical error made by the robot. Specifically, the frustration model of more patient users follow a near linear pattern, while impatient users follow an exponential pattern.

We combine interruption and error based components with equal weight, with an additional initialisation cost representing user's individual characteristics [32] $F = Init + IF + EF$ where $Init = 10 - w - p$.

Objective Function: The agent's goal is to perform actions $a \in A$ assisting their collaborating user $u \in U$ at various time slots $t \in T$, without inducing unnecessary frustration, causing them to exceed their individual frustration tolerance limit $\mathcal{B}$. As part of this, the agent has to estimate the oracle User Preference Tensor $\mathcal{M}$, while interacting with the user and update the estimated user preference tensor M based on the feedback received. We define an objective function identifying the optimal sequence of robot actions ($\mathcal{A}^*$):

$$\mathcal{A}^* = \arg\min_{\mathcal{A}} \frac{||\mathcal{M} - M(\mathcal{A})||_1}{|U| \times |T| \times |A|} \quad s.t. \ cost(\mathcal{A}) \leq \mathcal{B} \tag{4}$$

where $M(\mathcal{A})$ defines the user preference tensor following the execution of action sequence $\mathcal{A}$.

Solution Concept: Since we aim to approximate the oracle User Preference Tensor $\mathcal{M}$, which is unknown in prior to the interactions, the optimisation problem (Eq. 4) cannot be solved optimally, so we treat it as a constrained sequential decision-making problem.

Our framework supports multiple decision-making models, however, all variants follow the same base structure. The agent initialises the preference tensor estimate M with random values, then at each time slot an action $a \in A$ is chosen to be performed, following a decision-making models. After executing this action (either assistive or communication), the interacting user provides feedback, which the agent incorporates in its belief model.

Decision-making Models: The decision-making models belong to two broad categories: non frustration-aware and frustration-aware. We hypothesise that the frustration-aware model class will outperform the non frustration-aware class of methods considering user's perceived frustration.

As *non frustration-aware* baseline models, we use an example of commonly used reinforcement learning / multi-armed bandit style models, Upper Confidence Bounds (UCB) [38] and a model following the popular instruction-based approach to collaboration (that we refer to as "AlwaysAsk-Agent") used in teacher-student learning models.

We propose three decision-making approaches in the class of *frustration-aware models*.

1) Frustration-aware Greedy model: This model utilises the frustration estimation model to choose from the action space in a greedy manner. At each time slot, the model estimates the expected frustration cost induced by each potential action $a \in A$ incorporating the past interaction history, and accounting for the model's confidence in the action's correctness. After calculating the expected frustration cost of each action, the one with the lowest expected cost is selected to be executed. With each interaction, the preference tensor estimate M gets updated based on the received feedback, following the model update mechanism of a UCB agent.

2) Frustration-aware Knapsack model: The "FrustrationAwareKnapsackAgent" decision model aims to select actions that limit the frustration generated considering multiple future steps, treating it as a knapsack optimisation problem (for a detailed description of knapsack problems, see [39]). The weights associated with items (i.e. actions) are represented as frustration cost estimates, the values are defined as information gain and the knapsack's capacity is the user's frustration tolerance level. We calculate information gain from the difference between the parent entropy and average child entropy, motivated by information theory. After using a knapsack model to choose an action that maximises the information gain without inducing too much expected frustration on the user, that action is executed, and the user provided feedback is used to update the preference estimate tensor M following an UCB agent's update mechanism.

3) Frustration-aware UCB model: This approach ("FrustrationAwareUCBAgent") is built using a version of the Upper Confidence Bounds algorithm [38], designed for multi-armed bandit problems with budget constraints and variable action costs [40]. In our setting, the budget constraint is provided by the user's frustration tolerance limit and the variable costs are calculated by the frustration estimation model, with values calculated by the UCB component. In this model, matching our setting, both exploration and exploitation aspects are constrained by the budget.

V. Simulations

Our quantitative experiment aims to evaluate the models' performance in a number of experiments and test its robustness using user data from different interaction contexts (e.g. assistive manipulation, household assistance).

User Characteristics: In the simulations, users are assigned personality profiles, describing their behaviour characteristics displayed when interacting with the robot. Willingness (w) and patience (p) values and a frustration tolerance limit are randomly assigned, remaining constant throughout the experiment. Once a user's frustration level reaches their tolerance limit, they cease the collaboration.

User Behaviour Patterns: We simulate users' behaviour preferences using human behaviour datasets providing data on individuals' unique personal preferences for completing a specified task involving a fixed discrete number of time steps and a fixed set of actions, iterated over multiple iterations.

The *50 Salads Dataset* [41], captures 25 people preparing two salads each, suitably representing a manipulation and assistance context for an HRI scenario, in which individuals execute the same task following their own preferred way (as individuals take different steps when preparing the meal and use ingredients in different orders). The preference dataset contains 5 primary actions in the form of ingredients ('tomato', 'cheese', 'lettuce', 'dressing' and 'cucumber').

A collection of datasets produced by the Centre of Advanced Studies in Adaptive Systems (*CASAS*) as part of the

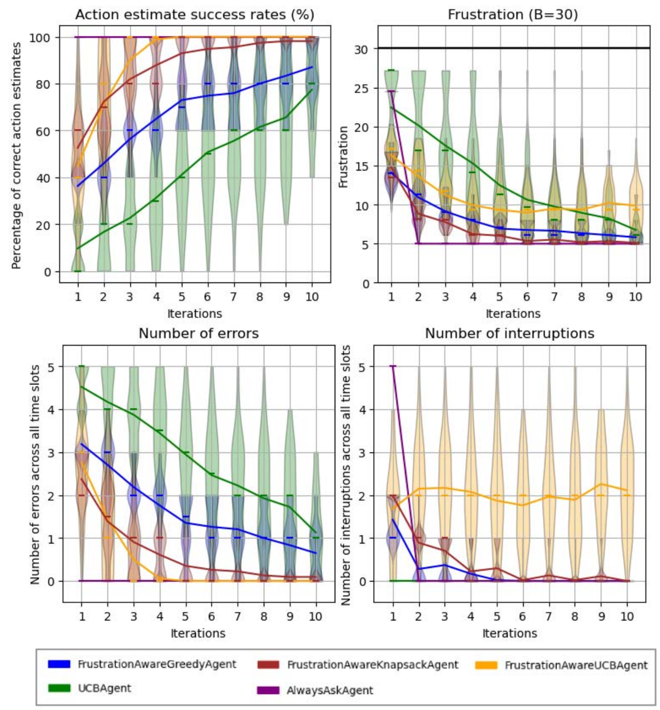

Fig. 2. Experiments ran using the 50 Salad dataset. Values show the mean of three experiment replicates, over all participating users, with the shaded areas representing the distribution of all measured values (with the shaded area's width being proportional to the number of measurements with at the given value). The second plot shows an example frustration tolerance level at the value 30 in black.

WSU CASAS smart home project represents user behaviour patterns and preferences in home environments. We used their dataset titled "WSU Smart Apartment Interweave ADL Testbed" [42] recording sensory events collected in a smart apartment testbed of the activities of 20 participants, who individually perform a fixed set of 8 household activities in the apartment, with everyone following their preferred order.

Evaluation Metrics: To evaluate the performance of individual models, we calculate the user preference estimation model's accuracy by using the objective function shown in Equation 4. We describe the model's decision-making behaviour by recording the number of logical errors and interruptions made and estimate the level of frustration throughout interactions using the previous frustration model.

Results: We mainly focus on the scenario with user behaviour patterns based on the 50 Salad dataset. The CASAS based simulations obtained similar results, however, for the sake of readability and presentation, these are omitted.

As shown in Figure 2, due to its continuous instruction queries, the AlwaysAskAgent yields the best preference estimation performance. All members of the frustration-aware class outperform the UCB model's performance, with FrustrationAwareUCBAgent reaching maximal success rates by iteration 4, and FrustrationAwareKnapsackAgent maintaining above 95% success rates from this time, converging towards 100%. FrustrationAwareGreedyAgent also shows continuously incremental performance with iterations, approximately doubling the mean success rate of the UCB model over all iterations.

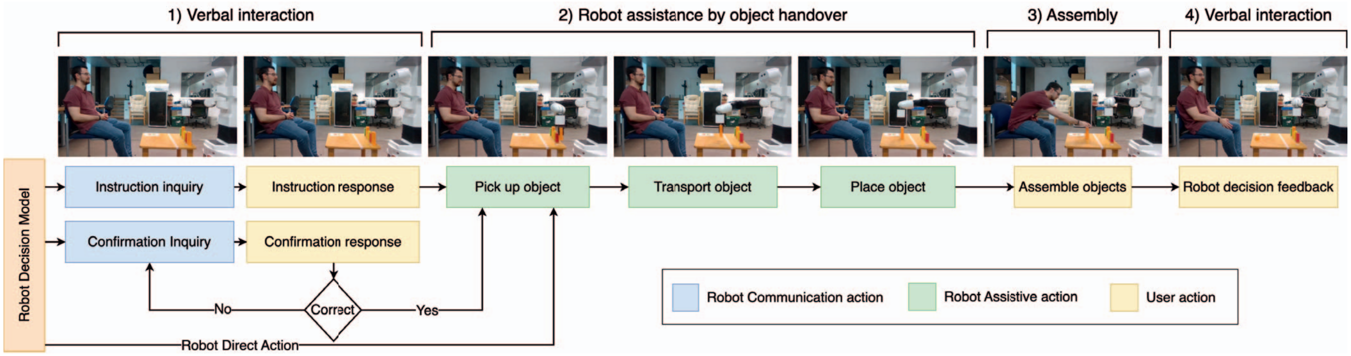

Fig. 3. Example of a participant's interaction with the robot in the burger assembly task. 1) Following on the robot's decision model, it selects the best way to elicit user's task preferences in the given situation (instruction inquiry, confirmation inquiry, direct action). 2) The assistive action is executed in the form of an object handover. 3) The object is received by the user, who performs the assembly and 4) provides feedback to the robot.

The FrustrationAwareUCBAgent maintains around 2 interruptions over time, due to the UCB-based strategy accounting for inconsistencies in user behaviour, resulting in the model's fast learning. The other two frustration-aware agents show similar tendencies to each other, timing their interruptions in the first few iterations, and rarely using interruptions later. The number of errors continuously decreases with all models, consistent with the preference estimation performance.

Resulting from its large number of interruptions, AlwaysAskAgent yields the highest initial frustration cost. The UCBAgent's generates slightly lower initial frustration levels, with near linear mean continuous decrease afterwards. The class of frustration-aware models outperform the others when considering their early frustration values by around 40% (13-17 versus 23-25 after the first iteration), and maintain continuous near exponential convergence towards the minimum, with FrustrationAwareKnapsackAgent converging the fastest. Due to their high initial frustration costs, non-frustration aware models may be suitable to fewer users (only those with high frustration tolerance levels).

VI. Human-Robot Interaction User Study

In our HRI setting, participants and the robot mimic a fast food restaurant collaborative burger assembly scenario. With ingredients (bread, meat, cheese, salad, bread) represented by coloured foam blocks, the robots performs a handover of the suitable object based on the preference model learnt from the human collaborator's responses and previous actions throughout the interactions. The human then completes the burger assembly by stacking the received object.

Experiment Procedure: Our experiment followed a repeated-measure design, each participant performed 9 trials (3 decision models (UCB, Instruction-based and FrustrationAwareGreedy), 3 repetitions) of the collaborative burger-assembly interaction, with interactions containing 5 action steps (i.e. burgers contain 5 ingredients). The order of models was randomised, not to affect the user's judgement. The participants were not informed of these conditions prior to the completion of the experiments.

We used the TIAGo Steel robot, equipped with a 7 DoF arm and parallel grippers. The robot base remained stationary during the interactions. The movement trajectories executing the handover are dynamically generated using the MoveIt framework [43], with additional user-safety constraints (mobile arm velocity, acceleration and accessible zones).

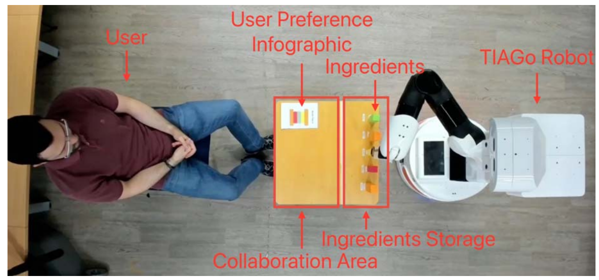

Fig. 4. The HRI setting from a top-down angle. During the collaboration, foam blocks ("ingredients") are initially placed on the robot's side of a table at fixed locations known to the robot to be picked up.

The natural language user responses to robot queries[3] are processed using Google's Speech-to-Text API with the audio feed recorded on a Logitech Yeti Nano microphone (using omnidirectional polar patterns) placed next to the experiment setup. User responses to confirmation requests or feedback requests are then pipelined into the Sentiment Analysis toolkit provided by Google's Natural Language API to binarily classify the user's response by its sentiment (i.e. if the user responded positively or negatively).

Participants: A total of 18 participants were recruited for the experiments. Data from 1 participant had to be discarded due to a mechanical failure of the robot. (The robot failed to close its parallel gripper due to a controller contact error, resulting in incomplete handovers. The issue was fixed after this trial.) Firstly, participants were provided with an Information Sheet and the experiment procedure was explained. Then a written informed consent was obtained, and each participant was assigned a numerical ID to anonymise the data. The whole experiment was conducted at the Active Laboratory at the University of Southampton, following experiment plans approved by the University's Ethics Committee.

Performance Metrics: Participants first completed a pre-interaction questionnaire collecting preliminary information: Frustration Discomfort Scale assessment [44], a Frustrative Nonreward Responsiveness Scale assessment [45] and a custom 5-point Likert-scale based section on users' prior experience with robots and their interaction expectations.

[3]Instruction inquiry such as "What should I do next?", action confirmation request as "Shall I pass you Ingredient 3?", or a feedback request following a robot decision such as "How did I do?".

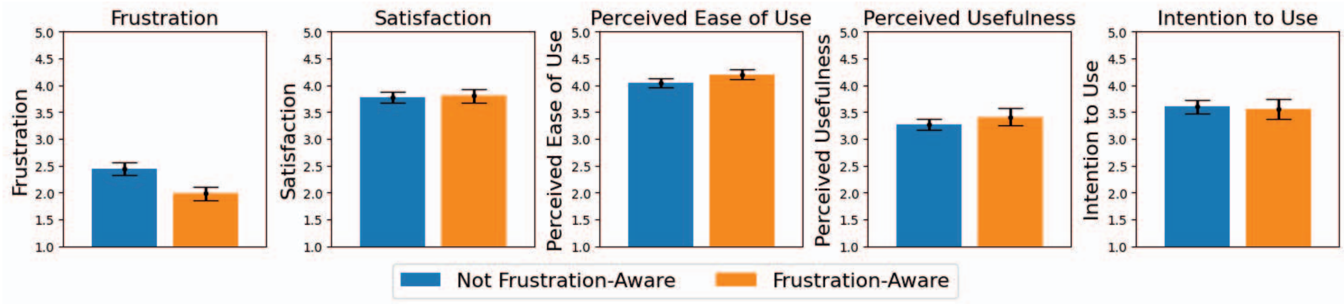

Fig. 5. Mean values with the error bars indicating standard error of the mean of user impressions of different behaviour models based on questionnaire outcome 1) Frustration, 2) Satisfaction with the Robot, 3) Perceived Ease of Use, 4) Perceived Usefulness, 5) Behavioural Intention to Use the Robot.

After each repetition, participants completed a post-interaction questionnaire. Data was collected following established prior work on Frustration [46], Satisfaction with the robot [47], Perceived Usability (using the Technology Acceptance Model or TAM) [48], [49], [50], Perceived Usefulness (TAM) and Behavioural Intention to Use (TAM), each using a 5-point Likert-scale. The models' quantitative performance was evaluated using Number of Interactions, (Instructions and Confirmation) and Completion Time metrics.

Fig. 6. Snapshot of multimodal footage captured throughout the user study, processed for the HRI dataset, with detected pose landmarks illustrated. 1) RGB side-view, 2) RGB top-down view, 3) RGB robot point-of-view, 4) depth robot point-of-view.

Multimodal Human-Robot Interaction Dataset: Synchronised multimodal recordings (see Figure 6) were made of each HRI burger-assembly scenario in the user study: 3 streams of RGB footage, 1 stream of depth footage and transcribed audio recordings of the interaction. User pose and face data is extracted from the RGB footage allowing for the future analysis of user behavioural responses to robot actions, while maintaining the anonymity of participants. Our dataset[4] contains over 6 hours of HRI footage, including various explicit and implicit user responses to robot actions.

Results: Evaluating the post-interaction questionnaire, the Mann-Whitney-Wilcoxon test was used to determine whether users provided different responses for the different models. We found significant differences ($\alpha = .05$) for the Frustration metric, where participants reported lower levels of frustration when interacting with the frustration-aware decision model ($p = .041$). The difference between models in other aspects of the questionnaire has shown similar trends (with users rating the frustration-aware model as more easy-to-use, useful and higher level of satisfaction with the robot), however, it did not reach the level of statistical significance.

[4]Publicly available at https://doi.org/10.5281/zenodo.14267885. For further details on the dataset's composition, see [51].

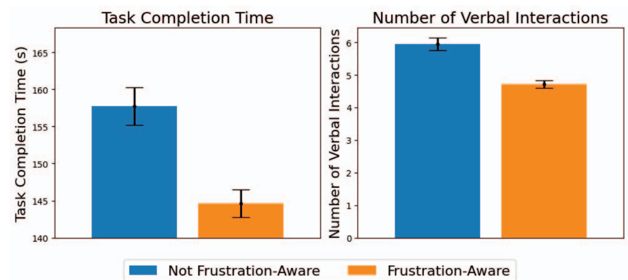

Fig. 7. Mean values with the error bars indicating standard error of the mean of 1) Task Completion Time, 2) Number of Verbal Interactions during the human-robot collaboration.

Participants welcomed when the frustration-aware model successfully learnt their preferences and only inquired when it was uncertain: "Robot is anticipating my needs.", "When it takes action without checking, it is always correct, so I trust it.". They welcomed the robot's suggestions (confirmation requests) when correct, "I liked how it suggested to pass me stuff on its own". However, they reacted negatively if the robot repeatedly offered incorrect suggestions (confirmation requests) or the interaction took longer, "I hope that the robot can respond more quickly to my instructions.".

Quantitatively, significant differences were found between models for the Number of Interactions ($p = .001$) and Completion time ($p = .001$), calculated using the Mann-Whitney-Wilcoxon test. Participants completed the collaboration in 9.04% less time using the Frustration-aware model and required 20.54% less number of interactions with the robot. Additionally, on average 38.77% of the interactions using the Frustration-aware model consisted of confirmation actions, requiring less cognitive effort from the participant.

VII. Conclusions

We proposed a Human-Robot Interaction (HRI) framework using Psychology-inspired frustration models, and combining multi-armed bandit and knapsack methods to address the issue of user frustration in collaborative HRI settings. We conducted simulations and a HRI user study using the TIAGo robot. Our findings show that users completed the collaboration task more efficiently using our frustration-aware model, while experiencing less frustration than with baseline models. There is still potential to improve user frustration modelling using fewer interactions. Additionally, we contributed a multimodal human-robot interaction dataset containing over 6 hours of implicit and explicit user responses.

References

[1] S. Elprama, I. El Makrini, B. Vanderborght, and A. Jacobs, "Acceptance of collaborative robots by factory workers: A pilot study on the role of social cues of anthropomorphic robots," in *The 25th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 2016, pp. 919–919.

[2] A. Meissner, A. Trübswetter, A. S. Conti-Kufner, and J. Schmidtler, "Friend or Foe? Understanding Assembly Workers' Acceptance of Human-robot Collaboration," *ACM Transactions on Human-Robot Interaction*, vol. 10, no. 1, pp. 1–30, 2021.

[3] Y. Yan and Y. Jia, "A Review on Human Comfort Factors, Measurements, and Improvements in Human–Robot Collaboration," *Sensors*, vol. 22, no. 19, p. 7431, 2022.

[4] A. Weidemann and N. Rußwinkel, "The Role of Frustration in Human–Robot Interaction – What Is Needed for a Successful Collaboration?" *Frontiers in Psychology*, vol. 12, 2021.

[5] W. Wang, Y. Chen, R. Li, and Y. Jia, "Learning and Comfort in Human–Robot Interaction: A Review," *Applied Sciences*, vol. 9, no. 23, p. 5152, 2019.

[6] K. Candon, Z. Hsu, Y. Kim, J. Chen, N. Tsoi, and M. Vázquez, "Perceptions of the helpfulness of unexpected agent assistance," in *Proceedings of the 10th International Conference on Human-Agent Interaction*, ser. HAI '22. Association for Computing Machinery, 2022, pp. 41–50.

[7] G. Bansal, B. Nushi, E. Kamar, E. Horvitz, and D. S. Weld, "Is the most accurate AI the best teammate? Optimizing AI for teamwork," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 13, pp. 11 405–11 414, 2021.

[8] J. Pagès, L. Marchionni, and F. Ferro, "TIAGo: The modular robot that adapts to different research needs," 2016.

[9] K. Bessière, J. E. Newhagen, J. P. Robinson, and B. Shneiderman, "A model for computer frustration: The role of instrumental and dispositional factors on incident, session, and post-session frustration and mood," *Computers in Human Behavior*, vol. 22, no. 6, pp. 941–961, 2006.

[10] A. Vinciarelli, M. Pantic, and H. Bourlard, "Social signal processing: Survey of an emerging domain," *Image and Vision Computing*, vol. 27, no. 12, pp. 1743–1759, 2009.

[11] N. Mirnig, G. Stollnberger, M. Miksch, S. Stadler, M. Giuliani, and M. Tscheligi, "To err is robot: How humans assess and act toward an erroneous social robot," *Frontiers in Robotics and AI*, vol. 4, 2017.

[12] 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)*, 2021, pp. 892–898.

[13] B. Nesset, D. A. Robb, J. Lopes, and H. Hastie, "Transparency in HRI: Trust and decision making in the face of robot errors," in *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '21 Companion. Association for Computing Machinery, 2021, pp. 313–317.

[14] M. Giuliani, N. Mirnig, G. Stollnberger, S. Stadler, R. Buchner, and M. Tscheligi, "Systematic analysis of video data from different human–robot interaction studies: A categorization of social signals during error situations," *Frontiers in Psychology*, vol. 6, 2015.

[15] M. Lewis, D. F. A. Hitchcock, and M. W. Sullivan, "Physiological and Emotional Reactivity to Learning and Frustration." *Infancy : the official journal of the International Society on Infant Studies*, vol. 6, no. 1, pp. 121–143, 2004.

[16] H. Jost, "Some Physiological Changes during Frustration," *Child Development*, vol. 12, no. 1, pp. 9–15, 1941.

[17] R. Hazlett, "Measurement of user frustration: A biologic approach," in *CHI '03 Extended Abstracts on Human Factors in Computing Systems*, ser. CHI EA '03. Association for Computing Machinery, 2003, pp. 734–735.

[18] J. F. Grafsgaard, J. B. Wiggins, K. E. Boyer, E. N. Wiebe, and J. C. Lester, "Automatically recognizing facial indicators of frustration: A learning-centric analysis," in *2013 Humaine Association Conference on Affective Computing and Intelligent Interaction*, 2013, pp. 159–165.

[19] J. F. Grafsgaard, J. B. Wiggins, A. K. Vail, K. E. Boyer, E. N. Wiebe, and J. C. Lester, "The additive value of multimodal features for predicting engagement, frustration, and learning during tutoring," in *Proceedings of the 16th International Conference on Multimodal Interaction*, ser. ICMI '14. Association for Computing Machinery, 2014, pp. 42–49.

[20] A. Kapoor, W. Burleson, and R. W. Picard, "Automatic prediction of frustration," *International Journal of Human-Computer Studies*, vol. 65, no. 8, pp. 724–736, 2007.

[21] L. Harrison, W. Dou, A. Lu, W. Ribarsky, and X. Wang, "Analysts aren't machines: Inferring frustration through visualization interaction," in *2011 IEEE Conference on Visual Analytics Science and Technology (VAST)*, 2011, pp. 279–280.

[22] H. A. Feild, J. Allan, and R. Jones, "Predicting searcher frustration," in *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR '10. Association for Computing Machinery, 2010, pp. 34–41.

[23] P. Trung, M. Giuliani, M. Miksch, G. Stollnberger, S. Stadler, N. Mirnig, and M. Tscheligi, "Head and shoulders: Automatic error detection in human-robot interaction," in *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, ser. ICMI '17. Association for Computing Machinery, 2017, pp. 181–188.

[24] D. Kontogiorgos, A. Pereira, B. Sahindal, S. van Waveren, and J. Gustafson, "Behavioural responses to robot conversational failures," in *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '20. Association for Computing Machinery, 2020, pp. 53–62.

[25] M. Stiber, R. Taylor, and C.-M. Huang, "Modeling human response to robot errors for timely error detection," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 676–683.

[26] Y. Mohamed, G. Ballardini, M. T. Parreira, S. Lemaignan, and I. Leite, "Automatic frustration detection using thermal imaging," in *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2022, pp. 451–459.

[27] Y. Cui, Q. Zhang, B. Knox, A. Allievi, P. Stone, and S. Niekum, "The EMPATHIC framework for task learning from implicit human feedback," in *Proceedings of the 2020 Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, J. Kober, F. Ramos, and C. Tomlin, Eds., vol. 155. PMLR, 2021, pp. 604–626.

[28] G. Li, H. Dibeklioğlu, S. Whiteson, and H. Hung, "Facial feedback for reinforcement learning: A case study and offline analysis using the TAMER framework," *Autonomous Agents and Multi-Agent Systems*, vol. 34, no. 1, p. 22, 2020.

[29] K. Candon, "Towards creating better interactive agents: Leveraging both implicit and explicit human feedback," in *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, 2023, pp. 3020–3022.

[30] N. Churamani, P. Barros, H. Gunes, and S. Wermter, "Affect-driven learning of robot behaviour for collaborative human-robot interactions," *Frontiers in Robotics and AI*, vol. 9, 2022.

[31] J. Dollard, N. E. Miller, L. W. Doob, O. H. Mowrer, and R. R. Sears, *Frustration and Aggression.*, ser. Frustration and Aggression. Yale University Press, 1939.

[32] M. Fleming and R. Cohen, "A decision procedure for autonomous agents to reason about interaction with humans," in *Proceedings of the AAAI 2004 Spring Symposium on Interaction between Humans and Autonomous Systems over Extended Operation*, 2004, pp. 81–86.

[33] R. Cohen, M. Y. Cheng, and M. W. Fleming, "Why bother about bother: Is it worth it to ask the user," in *Proceedings of AAAI Fall Symposium*, 2005.

[34] L. Berkowitz, "Frustration-aggression hypothesis: Examination and reformulation." *Psychological Bulletin*, vol. 106, no. 1, pp. 59–73, 1989.

[35] R. J. Thompson and R. H. Kolstoe, "Physical aggression as a function of strength of frustration and instrumentality of aggression," *Journal of Research in Personality*, vol. 7, no. 4, pp. 314–323, 1974.

[36] J. Breuer and M. Elson, "Frustration-aggression theory," in *The Wiley Handbook of Violence and Aggression*, P. Sturmey, Ed. Wiley Blackwell, 2017, pp. 1–12.

[37] R. J. Filer, "Frustration, satisfaction, and other factors affecting the attractiveness of goal objects." *The Journal of Abnormal and Social Psychology*, vol. 47, no. 2, pp. 203–212, 1952.

[38] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time Analysis of the Multiarmed Bandit Problem," *Machine Learning*, vol. 47, no. 2, pp. 235–256, 2002.

[39] H. Kellerer, U. Pferschy, and D. Pisinger, "Multidimensional Knapsack Problems," in *Knapsack Problems*. Springer, 2004, pp. 235–283.

[40] W. Ding, T. Qin, X.-D. Zhang, and T.-Y. Liu, "Multi-armed bandit with budget constraint and variable costs," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 27, no. 1, pp. 232–238, 2013.

[41] S. Stein and S. J. McKenna, "Combining embedded accelerometers with computer vision for recognizing food preparation activities," in *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ser. UbiComp '13. Association for Computing Machinery, 2013, pp. 729–738.

[42] G. Singla, D. J. Cook, and M. Schmitter-Edgecombe, "Tracking Activities in Complex Settings Using Smart Environment Technologies." *International journal of biosciences, psychiatry, and technology (IJBSPT)*, vol. 1, no. 1, pp. 25–35, 2009.
[43] D. Coleman, I. A. Sucan, S. Chitta, and N. Correll, "Reducing the barrier to entry of complex robotic software: a moveit! case study," vol. abs/1404.3785, 2014.
[44] N. Harrington, "The frustration discomfort scale: Development and psychometric properties," *Clinical Psychology & Psychotherapy*, vol. 12, no. 5, pp. 374–387, 2005.
[45] K. A. Wright, D. H. Lam, and R. G. Brown, "Reduced approach motivation following nonreward: Extension of the BIS/BAS scales." *Personality and Individual Differences*, vol. 47, no. 7, pp. 753–757, 2009.
[46] L. H. Peters, E. J. O'Connor, and C. J. Rudolf, "The behavioral and affective consequences of performance-relevant situational variables," *Organizational Behavior and Human Performance*, vol. 25, no. 1, pp. 79–96, 1980.
[47] P. A. Lasota and J. A. Shah, "Analyzing the effects of human-aware motion planning on close-proximity human-robot collaboration." *Human factors*, vol. 57, no. 1, pp. 21–33, 2015.
[48] V. Venkatesh and F. D. Davis, "A theoretical extension of the technology acceptance model: Four longitudinal field studies," *Management Science*, vol. 46, no. 2, pp. 186–204, 2000.
[49] F. D. Davis, "Perceived usefulness, perceived ease of use, and user acceptance of information technology," *MIS Quarterly*, vol. 13, no. 3, pp. 319–340, 1989.
[50] C. Bröhl, J. Nelles, C. Brandl, A. Mertens, and V. Nitsch, "Human–Robot Collaboration Acceptance Model: Development and Comparison for Germany, Japan, China and the USA," *International Journal of Social Robotics*, vol. 11, no. 5, pp. 709–726, Dec. 2019.
[51] B. Gucsi, N. T. V. Tuyen, B. Chu, D. Tarapore, and L. Tran-Thanh, "HRI-SENSE: A Multimodal Dataset on Social and Emotional Responses to Robot Behaviour," IEEE/ACM International Conference on Human-Robot Interaction (HRI), 2025, In Press.