2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Investigating the Role of Implicit Signals in Adaptive User-Aware Human-Robot Interactions

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

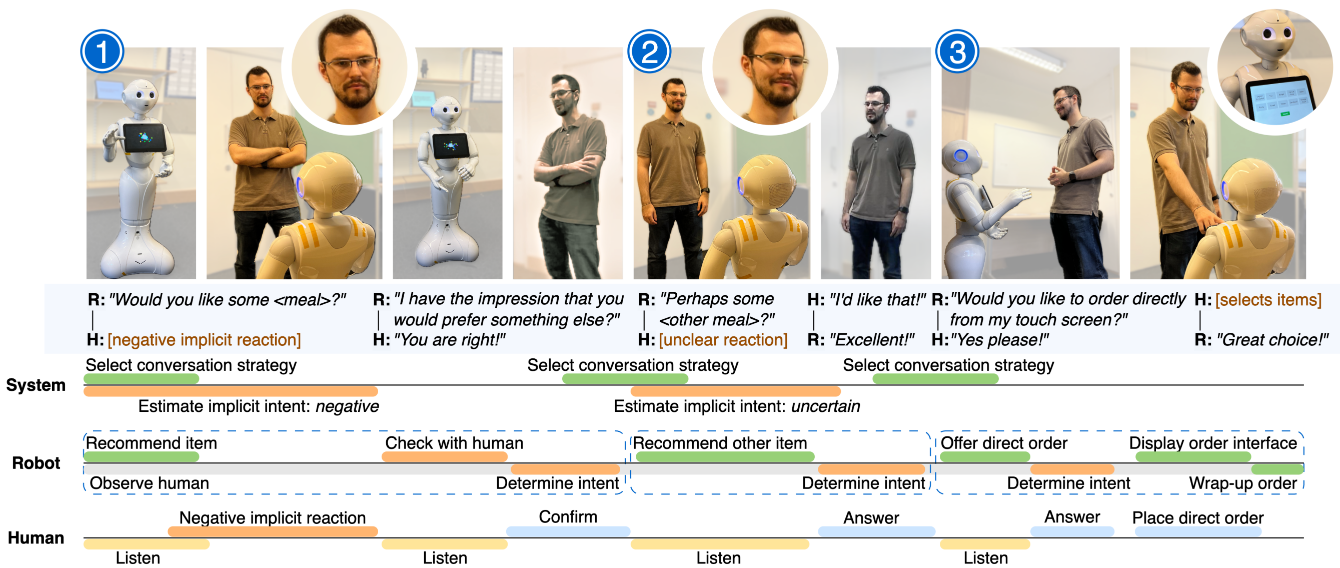

Fig. 1: We introduce an adaptive human-robot interaction framework building on hybrid (implicit and explicit) feedback signals designed to adjust the robot's behavioural style to match user preferences and personal characteristics without unnecessarily causing them frustration.

***Abstract*— Our work investigates how social robots can act in a user-aware manner by adapting their behaviour to users' personal characteristics and preferences without unnecessarily exposing them to frustration through the robot's actions. In particular, we investigate how implicit social signals inadvertently exhibited by users (e.g. facial expressions) during interactions can be incorporated into user-aware decision-making models while accounting for the systematic limitations of implicit feedback signals (e.g. inconsistency, noise, culture and individual-dependence). Doing so, we develop a user-aware adaptive decision-making and learning framework for human-robot interactions, building on implicit signal processing, cue-based intent inference, and multiarmed bandit learning techniques. Evaluating our approach, we conduct a user study where participants interact with a Pepper robot in a cafeteria style interaction scenario, with the robot providing recommendations and taking orders while adapting its behaviour to individual users. The experimental results demonstrate our proposed model's success in adapting its behaviour (i.e. conversational style) to users with different personal characteristics, while receiving 80% positive user feedback, and user questionnaire responses reporting higher perceived usefulness than baseline approaches. Questionnaire responses also illustrate positive user impressions of implicit signal based approaches while highlighting the importance of accounting for their limitations in learning models. In addition, we provide a dataset of over 5 hours of human and robot behaviour data extracted from multimodal recordings captured as part of our user study.**

I. Introduction

As the deployment of robots progresses from restricted factory-like environments to social settings such as households or public spaces, it becomes increasingly important to design robotic systems in a user-centric manner. With recent surveys [1], [2], [3] showcasing user acceptance and trust as the key factors to robotic systems' success in social settings, robot behaviour in human-robot interactions has to align with user preferences. Due to the difference between individuals' personal characteristics, tolerance levels and preferences, unlike in accuracy or efficiency centred unmanned robotic scenarios, "one-size-fits-all" robot behaviour approaches are insufficient for successful human-robot interactions.

In order for users to accept social robots' assistance and find them helpful [4], rather than sources of distraction or frustration, robots need to pursue user-aware behaviour: they need to learn to adapt to human factors including user preferences, values and affective state in addition to targeting objective performance metrics (e.g. model accuracy) [5].

In our work, we propose a user-aware decision making and learning framework for social robots, with a particular focus on investigating the utility of implicit social signals

[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].

[6] during human-robot interactions in everyday contexts. We address the questions: how can we incorporate implicit social signals such as facial expressions into social robot decision making in a user-aware manner with the robot inferring users' emotional and social state, learning to adapt its behaviour to individuals' characteristics, thus limiting their frustration; and how can the robot account for the potential limitations of implicit signals. Additionally, we address limiting the excessive frustration [7] caused by frequent robot errors (or non-preferred behaviours) encountered during the time-consuming and user-demanding training phase in typical HRI online learning models, through the use of domain-specific model pre-training and live few-shot user-adaptation.

During interactions, humans involuntarily exhibit implicit social signals [6] (e.g. facial expressions or gestures) indicative of their reaction to current events, social context or general affective state, making them a promising information source for user-aware robot behaviour adaptation. However, social signals are not always present and may differ at a cultural [8] or even individual level – limitations, which implicit signal based frameworks have to consider. Past research has explored social signals for error detection in HRI [9], adaptive learning applications [10], or teaching through implicit signals [11], [12], however, its incorporation to user-aware robotic decision making models is mostly unexplored. Particularly in "natural" scenarios where users do not encounter situations eliciting excessive social signals (e.g. critical robot error prompting surprised or frustrated expression). Additionally, participants interacting with the robot are not instructed to exaggerate emotions or use their expressions to teach the system unlike various previous works investigating implicit signals.

For the purpose of evaluating our user-aware learning framework, we assume a retail style cafeteria assistance interaction scenario (see Figure 1) with the robot providing suitable meal recommendations and taking orders while interacting with the user taking the customer role. The structured, yet open-ended design of this interaction scenario allows for robot behaviour to adapt to individual user characteristics through adjusting recommendations and interaction style. Further on, it enables continuous interactions and the exchange of multimodal (implicit and explicit) signals between the human-robot pair during the interaction.

Throughout the interaction, our framework learns how to adaptively adjust the robot's behaviour style to individuals' preferences building on a combination of implicit and explicit feedback signals used to infer users' affective state, frustration, intent and response to the robot's actions, while modelling the limitations of implicit signals using sequential learning multi-armed bandit methods [13], [14]. The proposed system is evaluated using quantitative and qualitative metrics against various baseline models as part of a user study conducted with 20 users that interacted with the Pepper robot [15] in our cafeteria-style scenario.

As an additional contribution, the HRI dataset composed of over 5 hours of anonymised human and robot behaviour data extracted from multimodal interaction footage recorded as part of the user study is publicly available.

II. Background

With humans typically following different behaviours depending on the social situation or interaction, they also exhibit a range of non-verbal signals (e.g. facial expressions [9], gestures[16], body posture[17], etc.) that differ between interaction contexts. As these involuntary signals require no cognitive effort from the individual, but are indicative of their social and affective state, they provide a suitable input for learning models addressing user characteristics, preferences, emotions or intentions in a proactive manner and thus are increasingly used in Social Signal Processing [6], Human-Computer Interaction (HCI) [18] and Human-Robot Interaction (HRI) contexts [19].

With approaches treating involuntary social signals as implicit user feedback, they have been successfully utilised for various inference purposes. User affective state information has been successfully extracted from social signals [10], with physiological signals (e.g. pulse or electrocardiogram (ECG)) providing the most reliable features [20]. However, due to its practical limitations, multimodal data-streams (e.g. gesture, gaze, pose) are often used instead in inference tasks for increased robustness in practical deployments. Other works utilised implicit features in smart tutoring applications [10] for detecting the presence of user frustration or predicting future frustration or in information retrieval systems [18].

Implicit social signals have shown promising results in the HRI domain as well, such as facial expression-based error detection during interactions [9], learning from implicit features when to solicit explicit user feedback [21], or inferring user impressions of current robot actions for robot task learning by interpreting implicit signals as a reward function [12] or evaluative feedback [11]. In addition to the previous works utilising social signals for inference or feedback in long-term teaching tasks, the work of [19] explores its use for robot behaviour adaptation as well. Their approach utilises speech input in conjunction with facial features for affective state inference with the aim of inferring the robot's intrinsic mood and selecting suitable actions in a negotiation context, but it does not consider the case of adaptive behaviour when relying solely on implicit feedback.

Other attempts for generating adaptive, user-aware robot behaviour in HRI settings build on Large Language Model (LLMs) architectures [22] including approaches for task personalisation using commonsense reasoning [23] or visual context-based conversation adaptation [24]. Additionally, approaches adapting robot behaviour in a user-aware manner utilising Psychological principle-based frustration inference models have also been proposed [25].

As the impact of implicit feedback for adaptive HRI has not been sufficiently addressed in previous works, we investigate how robots can adapt their behaviour in HRI scenarios to suit user characteristics and preferences by utilising implicit feedback features while accounting for its limitations, without exposing the user to excessive frustration during the interactions.

III. Problem Formulation

Multiple subproblems are addressed by the social robot during the cafeteria HRI scenario, while keeping user expe-

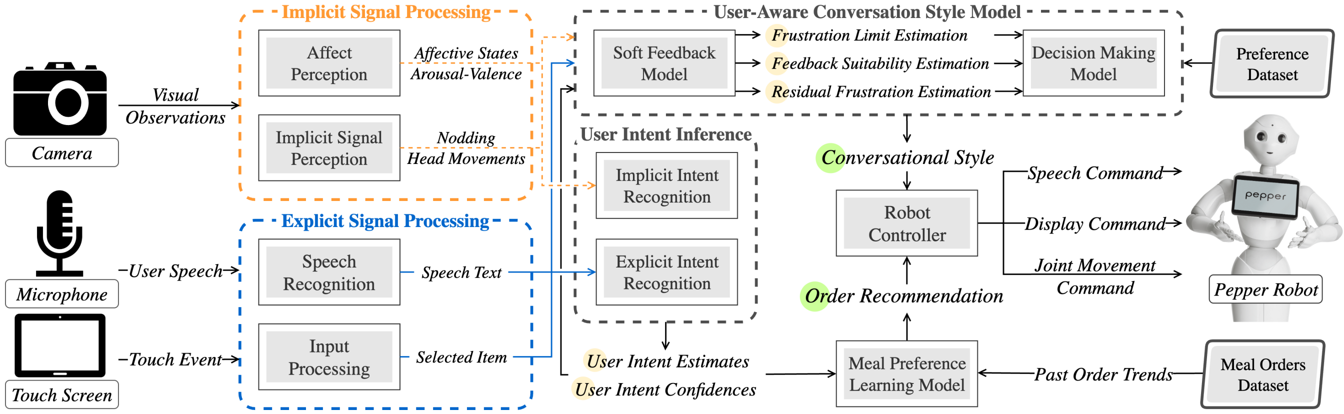

Fig. 2: Key components' interaction in the adaptive user-aware HRI framework designed for the utilisation of hybrid (implicit and explicit) input signals, deployed in the cafeteria-style HRI scenario. Implicit and explicit feedback signals used for user intent inference and adapting robot conversational behaviour to the user when offering suitable meal recommendations.

rience (i.e. individuals' frustration tolerance limit) in mind:
A) the recommendation subproblem: the robot has to learn what users' typical meal orders are, how these correlate with the time of the day and common meal combinations and provide suitable recommendations to customers.
B) the implicit inference subproblem: utilising social signals to understand users' implicit reactions (i.e. infer affective state and user intent) to robot actions.
C) the user-aware behaviour adaptation subproblem: the robot has to learn how to adapt its behaviour to best suit the user's individual characteristics, preferences and current affective state utilising implicit and explicit feedback to avoid unnecessarily frustrating them, while accounting for the limitations of implicit signals.

We formulate the overarching task as a sequential decision-making problem, with the aim of learning what actions the robot should take to match the user's individual preferences and characteristics, while constrained by the user's frustration tolerance level.

Doing so, our framework's (Figure 2) primary goal is to learn a policy π responsible for governing the robot's high-level behaviour by selecting an action $\mathbf{a}_n(t)$ for the robot to perform at time slot t that matches preferences of user n. As the robot executes action $\mathbf{a}_n(t)$, it may receive positive user feedback in the form of a reward $r(\mathbf{a}_n(t))$ if the action's value $v(\mathbf{a}_n(t))$ is below a user-specific threshold $\theta_n \in (0,1)$ (representing actions that match the user's personal preferences). We assume thresholds that are drawn independently from a distribution F, with both F and θ_n thresholds unknown to the robot. We consider reward functions to be identical for all users and a non-decreasing function of the action's value. If the action's value $v(\mathbf{a}_n(t))$ exceeds the threshold θ_n (representing actions that are unsuitable for user n, causing them frustration), their frustration tolerance budget B_n with initial value $\mathcal{B}$ is decreased by one unit. When users' frustration tolerance budget is exhausted, they abandon the interaction setting, resulting in a failed assistance attempt.

Accounting for the limitations of implicit feedback signals (e.g. not always exhibited, difficult to observe and interpret), a soft feedback model is applied. Specifically, feedback $\mathbf{1}(v(\mathbf{a}_n(t)) \leq \theta_n)$ is revealed to the robot (i.e. user feedback is exhibited, successfully observed and interpreted) for actions valued below the threshold θ_n (i.e. positive feedback for preferred robot actions) with probability $p_1 \in (0,1)$, and with probability $p_2 \in (0,1)$ for actions above the threshold (i.e. negative feedback for non-preferred robot actions), with p_1 and p_2 unknown to the robot. As such, the model accounts for implicit signals as partial probabilistic feedback. If explicit user feedback is provided, the robot observes that immediately (with $p_1 = 1$ and $p_2 = 1$).

IV. Adaptive User-Aware Human-Robot Interaction Framework

Action Space: With possible robot actions defined as $\mathbf{a} = (m, c)$, actions are composed of two components: meal recommendation m from the set of all available meal items M; and behaviour style c from all available styles C, regulating action execution. For practical examples, we assume conversational styles to represent behaviours in C, however, the framework is compatible with any behaviour.

Multiple robot behaviour styles are designed to cover a spectrum ranging from providing the user with lots of detailed recommendations, over limited recommendations, to offering to provide recommendations, and to offering a direct manual order option without any recommendations. With a gradually increasing amount of interaction and recommendations among styles, we assume each user to have a threshold (θ_n) under which they find the robot's help useful, but find it frustrating or unnecessary above.

Objective Function: The robot's goal is to perform an action $\mathbf{a}$ at each time t while adapting its behaviour to suit user n (regarding θ_n), without exceeding their individual frustration tolerance limit B_n. To select suitable actions, the robot simultaneously outputs an appropriate meal recommendation m given the time of the day and previous meal orders, updates learning models following user intents inferred from implicit and explicit feedback signals, and selects behaviour c to govern the interaction's manner. Doing so, it follows an objective function identifying the sequence of actions, that maximise the total discounted reward over all N users without exceeding users' frustration tolerance limit:

$$max_{\mathbf{A}_n \forall n \in N} \sum_{n} \sum_{t=1}^{T_n} \gamma^{t-1} r(\mathbf{a}_n(t)) \mathbf{1}(v(\mathbf{a}_n(t)) \leq \theta_n)$$
$$s.t.\ \forall n \in N : \left(\sum_{t=1}^{T_n} \mathbf{1}(v(\mathbf{a}_n(t)) > \theta_n) \right) \leq \mathcal{B} \tag{1}$$

where $\mathbf{A}_n$ denotes the sequence of actions interacting with user n, γ denotes a discount factor representing the diminishing effect of robot actions over time (i.e. the user's impression of the robot reflects recent actions more, with earlier actions gradually forgotten) and T_n denotes the total number of interactions with user n.

Solution Concept: Since users' preferences and characteristics (e.g. residual frustration (i.e. B_n), quality of exhibited feedback signals (i.e. p_1 and p_2) and robot behaviour tolerance levels (i.e. θ_n)) are unknown prior to the interactions, we address it as a frustration-constrained sequential decision-making problem. Solving this, multiarmed bandit techniques are used to output the behaviour component c of robot action $\mathbf{a}$ with an intermediary step of estimating p_1, p_2 and the distribution of typical user behaviour preferences F over time, while a separate recommendation preference learning model is responsible for the meal recommendation m component. For processing implicit and explicit user responses, facial feature-based affective state and intent inference models are used alongside the analysis of speech and touch input.

Meal Preference Learning Model: For providing meal recommendations to users that suit the time of the day and match previously ordered items, a Sequential k-Means Clustering model [26] with an additional domain-specific pre-training phase is used. This approach allows for the model to learn common patterns appearing in typical orders by identifying clusters in multidimensional past order datasets (i.e. order time and order component or inter order component correlations) and provide a suitable meal recommendation (e.g by selecting the item with the shortest distance in vector space from context data such as time or previous items in the current cluster). By utilising a pre-training phase on a domain-specific (meal orders) dataset, the model can learn typical user order behaviours offline without exposing users to unnecessary frustration during the training process, and adapt to the live user's meal preferences in a few-shot manner while simultaneously keeping the sequential k-Means model updated with new data.

Implicit Signal Processing: From the range of social signals exhibited by users during the interaction, our framework focuses on facial expressions and head movements since the cafeteria scenario is primarily stationary, making other bodily signals scarce or difficult to record. Facial signals are utilised in two manners: affect and implicit signal perception.

As part of affect perception a range of preprocessing steps involving face detection and cropping followed by discrete affective state classification is performed using the ENet architecture [27] trained on the VGGFace2 dataset [28] provided by the EmotiEffLib toolkit [29], [30] due to its robust real-time inference capabilities, outputting classification probabilities for 8 emotion classes. In addition, scalar arousal and valence values (in the (-1,1) range) positioning the encountered emotion on the spectrum of all affective states are also inferred at the rate of 3-5 frames per second, allowing for the rapid detection of changes in users' emotions.

Further implicit signals, specifically head movements indicative of gestures such as nodding behaviour were also extracted utilising a pipeline of face detection and alignment via MTCNN [31] and sliding-window based nodding gesture classification. Although head gestures are difficult to generalise due cultural variations [8], they have been successfully utilised in past works [32].

Explicit Signal Processing: Explicit feedback signals are collected from users in the form of direct speech in response to the robot's questions and touch button presses on the robot's tablet screen indicating the users' order preferences. Verbal input is processed using Google's Speech-to-Text API, while the result of touch inputs (i.e. menu items selected by the user) is directly forwarded.

User Intent Inference: Depending on the decision model used, features collected from both channels or explicit channels only are processed, aiming to extract user intentions.

In explicit feedback only scenarios, the user response transcripts are pipelined into a binary classifier (via the Sentiment Analysis toolkit by Google's Natural Language API) outputting the user response's sentiment (i.e. whether the user replied positively or negatively) and its intensity.

When having access to both feedback channels, implicit signals are usually exhibited naturally before speaking, thus these are addressed first. Given the scarcity of interaction datasets mapping implicit features to user intentions, the classification problem is approached in a cue-based model-free manner, successfully utilised in the context of empathy in previous works [33]. With affective state estimates, confidences, arousal-valence scores and head gesture classifications used as input, intent inference estimates along confidence values are output. If the confidence levels are below a threshold, the system falls back to explicit feedback.

User-Aware Behaviour (Conversation Style) Model: For learning to adapt robot behaviour to user preferences and characteristics utilising partially observable user feedback while accounting for user frustration limits, our model builds on a version of multiarmed bandit theory's Upper Confidence Bounds [13] algorithm designed for settings with budget constraints and partially observable feedback [14].

Since distribution F and values p_1 and p_2 are never revealed, the model works by first learning a noisy estimate of these using Linear Search Exploration on an exploration set of users (conducted in interaction with simulated users following realistic personality characteristics to avoid exposing live users to unnecessary frustration). Then a personalised strategy is performed on the live users following the estimates, aiming to infer individuals' residual frustration budget (B_n) in a UCB manner, by assigning an optimistic upper and lower confidence bound value to each potential action and continuously updating these bounds in an iterative manner while the action value estimates converge towards the true mean. Once a suitable behaviour is selected, it is used to adapt the execution of the recommendation action.

V. Human-Robot Interaction User Study

Our user study takes place in the context of a continuous contact retail assistance scenario (similarly to established HRI works on non-verbal behaviour [16]) providing structure, yet allowing for differences between interactions depending on user actions and robot behaviour. The scenario mimics a cafeteria ordering setting, where participants act as customers and Pepper provides meal recommendations based on a meal preference model learnt from past orders, while adapting its conversational style to match the user's preferences. The human responds to Pepper's suggestions and prompts verbally, while occasionally exhibiting involuntary implicit signals, with the robot utilising these features when learning to adapt its conversation style to suit the individual.

Experiment Setup: As part of the scenario participants stand opposite the Pepper robot, while the researcher monitors the interaction hidden from the participant's view (see Figure 3). This ensures that the participant fully focuses on Pepper without feeling inclined to consult the researcher in unexpected situations. The location is equipped with a recording setup consisting of two RGB cameras, an omnidirectional microphone and a tablet screen displaying information about the current experiment and the robot's speech transcript.

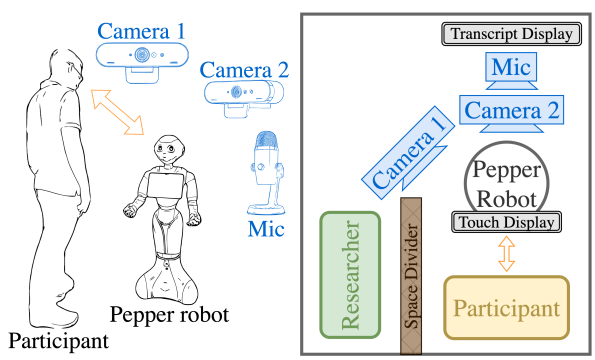

Fig. 3: The physical human-robot interaction setup designed for the user study and multimodal data collection illustrated.

Experiment Procedure: Prior to beginning the experiments, each user had a chance to get familiar with Pepper (e.g. when to respond, how to use touch screen) in a separate introductory scenario consisting of small talk. This component has been solely designed to mitigate the novelty effect when users encounter behaviour models in the main scenario, without being used for other purposes.

Our experiment is conduced in a repeated-measure, randomised order within-subject fashion, with participants interacting with Pepper throughout 6 trials (3 robot behaviour models (*Explicit-signal Touch-based (ET)*, *Explicit-signal Conversational (EC)*, *Hybrid-signal Conversational User-Aware (HC-UA)*), 2 interactions each time). The order of robot behaviour models is randomised between participants to avoid bias. Each interaction contains up to 10 steps (e.g. robot recommendations) depending on the behaviour model and the user's reactions and feedback. Ensuring the naturalness of implicit user reactions in a context where strong emotions are rarely elicited, participants were not asked to use expressive or exaggerated gestures/mimics or attempt to convey information to teach Pepper implicitly.

During interactions, Pepper's base remained stationary, while its speech was complemented with animated speech gestures (i.e. suitable arm and head movements). All behaviour modes animated speech with gestures in the same manner. Ensuring participants safety, Pepper paused its movements when accepting user input (e.g. using the tablet) and implemented additional user safety constraints such as velocity and acceleration limits.

Pepper accepted input from users in the form of natural language verbal communication and via buttons displayed on its tablet screen (indicating meal options). To avoid unnecessary frustration caused by communication misunderstandings due to users speaking too early (i.e. before Pepper started listening), Pepper was programmed to indicate when it is ready to receive verbal user input with green colours on the external tablet screen and LED lights on Pepper's ears. The framework handling input feature processing and running decision making and learning models was run on an external computer, exchanging information with Pepper over the local network via the ZeroMQ protocol.

Domain-specific Training Data: Assisting the frustration-free pre-training of recommendation and behaviour style learning models, datasets with typical meal order and user behaviour preferences were utilised in simulation. Specifically, a dataset containing 1000 orders with corresponding timestamps was generated following typical meal order patterns (e.g. meal pairings, meal time trends), simplifying the selection of established datasets [34], [35] to 10 available menu items for the ease of user interaction. Further on, user preferences on the scale of available robot behaviours were initially modelled as a uniform $U[0, 1]$ distribution function F, while the exhibited implicit feedback quality values (p_1, p_2) were modelled based on typical AffectNet [36] 8-class classification model benchmarks.

Participants: Overall 20 participants of different ages ($M = 28, SD = 8.26$), medium amount of technological/robotics experience ($M = 3.29, SD = 1.12$ recorded on a 5-point Likert scale with 5 being very experienced) coming from various cultural backgrounds interacted with the Pepper social robot as part of the user study. Prior to the experiments, individuals were provided with the study details, then a written informed consent was obtained. Each participant was assigned a numerical ID to anonymise their data. The study was conducted at the University of Southampton, with experiment plans approved by the University's Ethics Board.

Benchmarks: As part of the user study, the performance of our proposed framework building on a combination of implicit and explicit feedback signals while dynamically adapting robot behaviour to suit the user (denoted as *HC-UA*) is evaluated against two baselines: *EC* and *ET* models.

While *HC-UA* utilises primarily implicit, secondarily explicit signals and adapts its behaviour to the user, *EC* solely utilises explicit feedback signals and follows a single

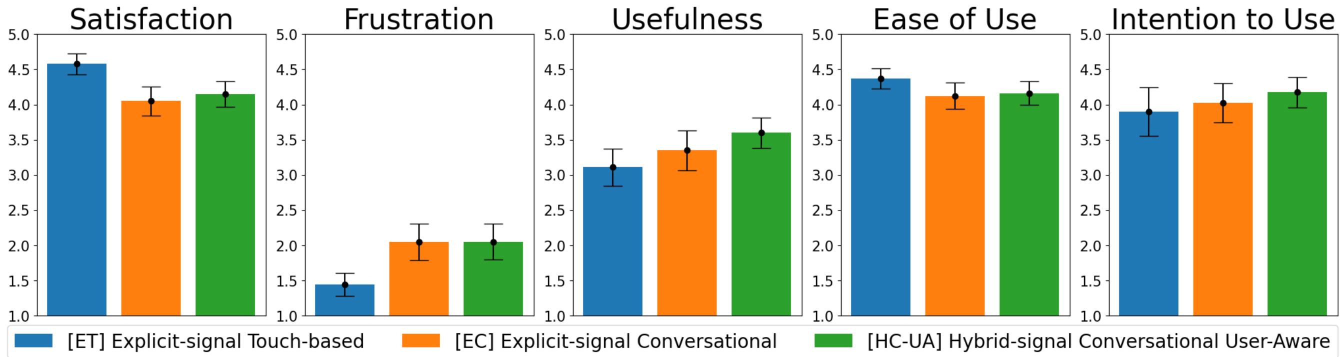

Fig. 4: User impressions of all robot behaviour models captured as part of the user study with data on 1) Satisfaction with the Robot, 2) Frustration, 3) Perceived Usefulness, 4) Perceived Ease of Use and 5) Behavioural Intention to Use the Robot. Bar charts showcase the mean results with error bars indicating the confidence interval (1.96*SE).

conversational behaviour while interacting with the user without adapting its behaviour. Similarly, *ET* builds on explicit feedback only, but interacts with users by taking orders directly using its tablet display without conversing or offering recommendations.

Measures: Preceding the human-robot interactions, user input was collected using a 5-point Likert scale questionnaire on their past experience with robots, interaction expectations, general background, personality via BFI–S [37] and their attitude towards handling frustration via the Frustration Discomfort Scale [38]. After interacting with each Pepper behaviour model, participants were asked to complete a post-interaction questionnaire on their impressions of robot behaviour. Data collection followed established work on Frustration [39], Satisfaction with the robot [40], Perceived Usability, Perceived Usefulness and Behavioural Intention to Use (using the Technology Acceptance Model) [41], [42], using a 5-point Likert scale. Additionally, explicit and implicit user feedback was collected during each interaction. For quantitative model evaluation, Number of Interaction Events (Touch Inputs, Recommendations Provided) and Implicit Inference Success metrics were used.

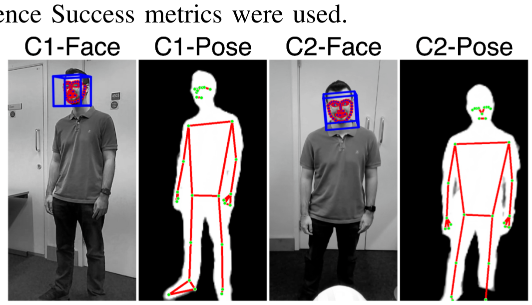

Fig. 5: Multimodal data streams (RGB video, audio) captured of the user during the human-robot interaction as part of the user study with extracted facial landmark, facial action unit, head pose and pose landmark data illustrated.

Multimodal Human-Robot Interaction Dataset: As part of our user study, a synchronised multimodal data recording was made of each human-robot interaction showcasing various implicit and explicit user reactions to robot behaviour. Recordings consist of two RGB camera feeds, audio recordings of interactions (transcribed) and robot action logs. User pose landmarks, head poses, facial landmarks and facial action units were extracted from the RGB footage (using OpenFace [43] and the MediaPipe Pose Landmarker [44] tools) and interaction transcripts from the audio data (using the Whisper large sized transcription model [45]), enabling the characterisation of user behaviour without user privacy concerns. Our multimodal dataset[1] captures 5 hours of HRI footage, composed of 240 interactions with 20 users.

VI. Results and Discussion

Overall User Impressions: Post-interaction questionnaire responses were analysed using the Independent-Samples Kruskal-Wallis test (chosen due to the collected data's deviation from normality assumptions) to determine whether users' self-assessed impressions differ for different robot behaviour models. Addressing the provided qualitative metrics over all samples (see Figure 4), although in cases of Satisfaction, Frustration and Ease of Use, the *ET* model outperformed the conversational models (with *HC-UA* slightly outperforming or being on par with *EC*), considering Usefulness and Intention to Use *HC-UA* yielded the best results, showing significant differences ($\alpha = .05$) for Usefulness.

The positive results for *ET* (regarding Satisfaction, Frustration and Ease of Use) can be explained by users' familiarity with the touch-based order interfaces widely available in fast food restaurants. Notably, however, users preferred it more regarding Usefulness and Intention to Use when the robot's behaviour enabled them to engage in conversation with the robot and receive recommendations, particularly with *HC-UA*. Also emphasised by the qualitative open-ended feedback received from users such as "enjoyed having the option to interact further and be recommended options", with others finding the touch-based *ET* approach limiting despite its efficiency "no different to smartphone app".

Role of Implicit Social Signals: Focusing on the role of implicit social signals in the human-robot interactions (Figure 6), results (column 1) show that implicit inference often cannot detect user intention with enough confidence to act on the information with ~75% of cases yielding uncertain outcomes. This may be due to the problem's difficulty: not only do users' facial and head gestures differ by culture and on an individual level, assuming successful feature

[1]Publicly available at https://doi.org/10.5281/zenodo.17121105.

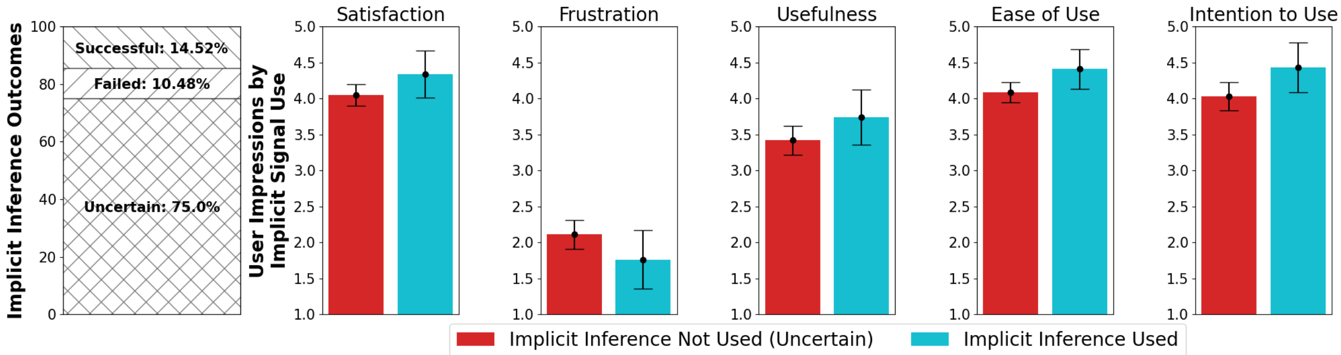

Fig. 6: Showcasing the role of implicit feedback signals in our HRI behaviour models. Outcomes of implicit inference attempts illustrated in column 1. Remaining columns illustrate user impressions of robot behaviours grouped by cases when [Red] implicit inference did not yield results with enough confidence to utilise in robot decision making, and [Blue] when inference from implicit signals yielded suitable results to use in robot decision making.

identification (e.g. one person's expression of happiness may correspond with another's surprise), but the context is also highly noisy during "in-the-wild" settings (e.g. hand covering mouth, long hair blocking eyes, glasses reflecting light).

However, when implicit features were suitable for inference, the models inferred intention accurately in $\sim$60% of cases. Importantly, when comparing user's impressions between cases where implicit inference yields enough confidence for the robot to act upon information from social signals and where it does not (see columns 2-6 in Figure 6), users preferred the former behaviour utilising implicit signals in all five metrics, with the Mann-Whitney-Wilcoxon test ($\alpha = .05$) yielding significant differences for Ease of Use ($p = .05$). Qualitative participant comments also indicate positive impressions on social signal inference in *HC-UA* "Robot understood when I did not want something before I spoke, removed awkwardness of rejecting recommendation.".

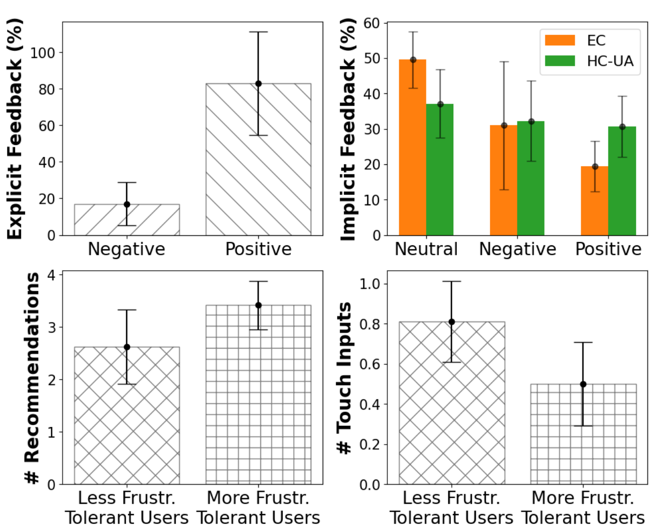

Fig. 7: Illustrating qualitative user impressions and quantitative results on *[HC-UA] Hybrid-signal Conversational User-Aware* model of our proposed framework. Top row showcases qualitative data based on explicit and implicit user feedback on *HC-UA*. Bottom row displays quantitatively how the robot (following *HC-UA*) adjusts its behaviour accommodating different user characteristics.

Conversational Adaptive Behaviour: The previous findings' support of the hypothesis that the usability of implicit social signals is limited in everyday scenarios (without extreme events causing intense reactions) highlights the importance of accounting for this uncertainty in behaviour models, as proposed in our *HC-UA* model.

Qualitative results collected on *HC-UA* during interactions (see Figure 7 top row) show primarily ($\sim$80%) positive explicit user feedbacks, while implicit social signals collected during the interactions also indicate more positive responses for *HC-UA* in comparison with the *EC* baseline.

Quantitative metrics (see Figure 7 bottom row) characterising robot behaviour (i.e. Number of Interaction Events) demonstrate that *HC-UA* offered significantly more recommendations (Mann-Whitney-Wilcoxon test with $p = .026$) to participants with higher self-reported frustration tolerance (unknown to the robot a-priori) then to those with lower self-reported frustration tolerance. Similarly, less frustration tolerant users were offered the less cognitively intensive direct touch input ordering opportunity by *HC-UA* significantly more times ($p = .048$), thus successfully adapting robot behaviour to individuals' preferences and characteristics.

VII. Conclusions and Future Work

We proposed a user-aware adaptive Human-Robot Interaction framework addressing the use of implicit social signals in HRI settings. We utilised multiarmed bandit learning, affective state estimation and cue-based intent inference techniques enabling adaptive robot behaviour in a frustration-aware manner, while accounting for implicit feedback signals' limitations at a decision making level. A user study using the Pepper social robot for evaluating our approach has demonstrated that participants found our *HC-UA* model utilising implicit signals to adapt its behaviour in a user-aware manner more useful than baselines, despite the touch-based interaction model (*ET*) being more efficient. Questionnaire results have shown that despite various systematic constraints limiting the usability of implicit signals, users prefer robots that utilise them, highlighting the importance of accounting for these constraints in decision making. Finally, we found qualitative and quantitative evidence in support of our proposed *HC-UA* model successfully adapting its behaviour to users, resulting in positive impressions. We also contributed a multimodal dataset of over 5 hours of human-robot interactions with key user features extracted.

As part of future works, various different multimodal implicit inference models and more flexible recommendation models, possibly based on LLMs are to be explored, in conjunction with evaluating the performance of our proposed framework in different interaction scenarios.

References

[1] 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.

[2] A. Meissner, A. Trübswetter, A. S. Conti-Kufner, and J. Schmidtler, "Friend or Foe? Understanding Assembly Workers' Acceptance of Human-robot Collaboration," *ACM THRI*, 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] K. Candon, Z. Hsu, Y. Kim, J. Chen, N. Tsoi, and M. Vázquez, "Perceptions of the helpfulness of unexpected agent assistance," in *ACM HAI*, 2022, pp. 41–50.

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

[6] 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.

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

[8] S. Kita, "Cross-cultural variation of speech-accompanying gesture: A review," *Language and Cognitive Processes*, vol. 24, no. 2, pp. 145–167, 2009.

[9] M. Stiber, R. H. Taylor, and C.-M. Huang, "On using social signals to enable flexible error-aware HRI," in *ACM/IEEE HRI*, 2023, pp. 222–230.

[10] 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 *ACM ICMI*, 2014, pp. 42–49.

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

[12] Y. Cui, Q. Zhang, B. Knox, A. Allievi, P. Stone, and S. Niekum, "The EMPATHIC framework for task learning from implicit human feedback," in *PMLR CoRL*, vol. 155, 2021, pp. 604–626.

[13] 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.

[14] A. Rangi, M. Franceschetti, and L. Tran-Thanh, "Sequential choice bandits with feedback for personalizing users' experience," 2021. [Online]. Available: arXiv:2101.01572

[15] A. K. Pandey and R. Gelin, "A mass-produced sociable humanoid robot: Pepper: The first machine of its kind," *IEEE Robot. Autom. Mag.*, vol. 25, no. 3, pp. 40–48, 2018.

[16] N. T. V. Tuyen, M. Adamski, Y. Wang, S. Okazaki, and O. Celiktutan, "The impact of robot's body language on customer experience: An analysis in a cafe setting," in *ACM/IEEE HRI*, 2023, pp. 248–252.

[17] A. Kleinsmith and N. Bianchi-Berthouze, "Affective body expression perception and recognition: A survey," *IEEE ITAC*, vol. 4, no. 1, pp. 15–33, 2013.

[18] H. A. Feild, J. Allan, and R. Jones, "Predicting searcher frustration," in *ACM SIGIR*, 2010, pp. 34–41.

[19] 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.

[20] M. Lewis, D. F. A. Hitchcock, and M. W. Sullivan, "Physiological and Emotional Reactivity to Learning and Frustration." *Infancy*, vol. 6, no. 1, pp. 121–143, 2004.

[21] K. Candon, "Towards creating better interactive agents: Leveraging both implicit and explicit human feedback," in *AAMAS*, 2023, pp. 3020–3022.

[22] M. A. Uddin and H. Salam, "Beyond rule-based context awareness: Large language models as adaptive cognitive layers in cyber-physical systems," *AAAI Symposium Series*, vol. 6, no. 1, pp. 140–147, 2025.

[23] M. Patel and S. Chernova, "Robot Behavior Personalization from Sparse User Feedback," in *IEEE RO-MAN*, 2025.

[24] R. Janssens, P. Wolfert, T. Demeester, and T. Belpaeme, "Integrating visual context into language models for situated social conversation starters," *IEEE ITAC*, vol. 16, no. 1, pp. 223–236, 2025.

[25] B. Gucsi, N. T. Viet Tuyen, B. Chu, D. Tarapore, and L. Tran-Thanh, "User-aware collaborative learning in human-robot interactions," in *IEEE ICRA*, 2025, pp. 7399–7406.

[26] V. Guralnik and G. Karypis, "A scalable algorithm for clustering sequential data," in *Proceedings 2001 IEEE International Conference on Data Mining*, 2001, pp. 179–186.

[27] A. Paszke, A. Chaurasia, S. Kim, and E. Culurciello, "Enet: A deep neural network architecture for real-time semantic segmentation," 2016. [Online]. Available: arXiv:2204.02681

[28] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "VG-GFace2: A dataset for recognising faces across pose and age," in *International Conference on Automatic Face and Gesture Recognition*, 2018.

[29] A. Savchenko, "Facial expression recognition with adaptive frame rate based on multiple testing correction," in *PMLR ICML*, A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, Eds., vol. 202, 2023, pp. 30 119–30 129.

[30] A. V. Savchenko, L. V. Savchenko, and I. Makarov, "Classifying emotions and engagement in online learning based on a single facial expression recognition neural network," *IEEE ITAC*, 2022.

[31] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, 2016.

[32] A. Samanta and T. Guha, "Emotion sensing from head motion capture," *IEEE Sensors Journal*, vol. 21, no. 4, pp. 5035–5043, 2021.

[33] C. Arzate Cruz, E. C. Montiel-Vázquez, C. Maeda, and R. Gomez, "When and how to express empathy in human-robot interaction scenarios," in *IEEE RO-MAN*, 2025.

[34] A. Pambudi, "Analyzing International Restaurant Orders Dataset," Kaggle, 2023. [Online]. Available: https://www.kaggle.com/datasets/agungpambudi/analyzing-restaurant-orders-international-dataset

[35] "Takeaway Food Orders - updated," Kaggle, 2019. [Online]. Available: https://www.kaggle.com/datasets/henslersoftware/19560-indian-takeaway-orders

[36] A. Mollahosseini, B. Hasani, and M. H. Mahoor, "AffectNet: A database for facial expression, valence, and arousal computing in the wild," *IEEE ITAC*, vol. 10, no. 1, pp. 18–31, 2019.

[37] F. R. Lang, D. John, O. Lüdtke, J. Schupp, and G. G. Wagner, "Short assessment of the Big Five: Robust across survey methods except telephone interviewing," *Behavior Research Methods*, vol. 43, no. 2, pp. 548–567, 2011.

[38] N. Harrington, "The frustration discomfort scale: Development and psychometric properties," *Clinical Psychology & Psychotherapy*, vol. 12, no. 5, pp. 374–387, 2005.

[39] 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.

[40] 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.

[41] 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.

[42] 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.

[43] T. Baltrušaitis, P. Robinson, and L.-P. Morency, "OpenFace: An open source facial behavior analysis toolkit," in *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2016, pp. 1–10.

[44] C. Lugaresi, J. Tang, H. Nash, C. McClanahan, E. Uboweja, M. Hays, F. Zhang, C.-L. Chang, M. G. Yong, J. Lee, W.-T. Chang, W. Hua, M. Georg, and M. Grundmann, "MediaPipe: A framework for building perception pipelines," 2019.

[45] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. Mcleavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *PMLR ICML*, A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, Eds., vol. 202, 2023, pp. 28 492–28 518.