RoboButler: Frustration-Aware Assistive User Localisation for Social Robots in Office Environments

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

Abstract— **In human-robot interactions (HRI), it is crucial for robots to be accepted by users and that they find robotic assistance attempts helpful rather than frustrating. Working towards this goal, we investigate the problem of frustration-aware robot behaviour planning in human-robot interaction contexts without continuous user contact or live feedback. Specifically, we address the question of how social robots can efficiently localise users and assist them with errands of various importance in office environments, while minimizing the frustration experienced by their human colleagues to enhance the overall interaction experience. Doing so, we design a frustration-aware decision-making and learning framework building on multiarmed bandit approaches and knapsack algorithms, in addition to developing a Psychology-based model of frustration tailored for HRI settings with limited user contact. Then we evaluate our approach on realistic user behaviour datasets, simulating the interactions' robotic components in Gazebo with a TIAGo robot, and perform further scalability analysis in graph-based simulations. The experimental results demonstrate that the proposed framework achieves localisation success rates and travel times that converge towards oracle values (outperforming other structured learning benchmarks) while yielding an estimated up to 75% less frustration – indicating the proposed framework's suitability for advancing to user studies and deployment in real-world scenarios.**

I. Introduction

Intelligent social robots [1], [2] are deployed in everyday environments seeking ways to provide intelligent systems that interact with or assist human users in their daily tasks, ranging from engaging in conversations, delivering coffee at the right time [3], preparing packages [4], or delivering laboratory samples and helping elderly people [5]. However, despite technological advances in recent years, the adaptation of such systems to social settings remains limited, with recent studies [6], [7], [8] highlighting the user's acceptance of robots as the key. The impressions a user formulates of robotic behaviour during an interaction are crucial to this [9], [10]: robot actions considered similar through objective metrics (e.g. efficiency, speed of task completion) may be perceived as helpful, successful; or constraining, frustrating, depending on the way they are executed (e.g. interaction timing or considering individual preferences). Thus, in order to develop human-robot interaction models that users actually find helpful and accept rather than find frustrating and hindering, our work focuses on modelling user impressions, frustration and user-aware behaviour models.

In particular, in this study we address interaction scenarios without continuous user feedback where human-robot contact only occurs periodically. Such interactions are typical in non-industrial, everyday environments facilitated by household or Office Assistant Robots (OARs). As platforms serving as conversation partners [11], running simple but time-consuming or repetitive errands [12], [13] they interact with various members of an office for limited intervals and often user feedback can only be provided retrospectively.

Our work investigates a setting where multiple office workers request an OAR's assistance with different navigation-based errands of various importance at predetermined times during the day tailored to individuals' needs. For instance: deliver a sandwich to Alice at 12:30, transport a small parcel from reception to Bob at 10:00, bring medicine and a glass of water to Charlie at 12:00. Given the nature of the assistive tasks, we refer to our interaction scenario as RoboButler as our robot has to strive for good performance with maintaining a tactful (non-frustrating) manner.

In most interactions, it is assumed that the user being assisted is near the social robot [14], or their location is known. This may not hold in real-world scenarios such as with OAR interactions since the robot interacts with multiple users throughout the day, who freely move around the environment following their own schedule, which the robot may not have access to. As such, in order for the OAR to assist, the user has to be located first. To do so, while navigating in the office, the OAR actively explores the environment identifying nearby users, potentially engaging in interactions, just like a human colleague would do.

However, if the robot spends too long exploring the environment, interacting with other users, or fails to locate the user requiring assistance in time, the user will experience frustration and dissatisfaction with the robot's performance. Additionally, if the robot's actions (either navigation or manipulation related when within sight) are not in agreement with the user's expectations, the user's frustration level might further increase. As a result, the user's disappointment and frustration may cause them to give up on expecting the robot's assistance. Moreover, they might lose interest in interacting with the robot in the future, as well [8], [15].

Unlike typical *"Where is Waldo?"* style user localisation problems [16] where the models aim to minimise the localisation time of *"Waldo"*, in our setting, due factors such as task importance, user tolerance or previously encountered users, the shortest localisation time may not be ideal. E.g.

[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]. For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) license to any Author Accepted Manuscript version arising. This paper has been accepted by the 2025 34th IEEE International Conference on Robot and Human Interactive Communication (ROMAN).

a longer localisation time for less important tasks might be equally acceptable as a faster time for important tasks or impatient users. Additionally, individuals' way of handling frustration (and tolerance limits) differ and affect how one reacts to localisation strategies that might be considered identical in *"Where is Waldo?"* localisation settings [16].

In this study, we address the questions of what would be an efficient way of selecting a path to explore to localise the user without causing them frustration? Or, what moves may a robot make to maximise the probability of exploring a sufficient area to localise the user before causing the user inconvenience? When should the robot explore unknown locations, and when should the robot visit locations frequented by the user? To address these questions, we design a user localisation framework tailored for social robots, considering potential user frustration. The robot formulates the user localisation process as a tensor estimation problem using heuristics from multi-armed bandit theory for updating values and knapsack optimisation to choose which path to explore. A path consists of a set of locations connected by a set of moves. The robot has to stop exploring once the user's frustration limit has been reached – after which the robot's assistance is no longer useful to the user.

To demonstrate and evaluate the proposed model, we design a simulator based multi-room office environment for a TIAGo Steel collaborative robot [17], and develop a customisable graph-based abstract environment simulator to efficiently run a large number of experiments in numerous different environment configurations. We model user behaviour in the office following established user behaviour datasets [18], [19]. The experimental results show that using our heuristics, we can localise users without causing them excessive frustration in multi-room environments of various sizes and configurations with near-optimal success rates in under 30% of the time required to traverse the whole office environment.

II. Related Work

Studies on User Location Discovery (ULD) [20] or Indoor Localisation Systems (ILS) [21], typically utilise pre-installed fixed sensor networks [22], [23] in the operational area (e.g. cameras [24], PIR [25]) resulting in coverage limitations to a predetermined area, extensive installation costs and no adaptation capabilities. This constraint makes them unideal for interactions where the operational area may change at short notice (such as a robot being deployed in different office levels) or where the deployment's budget is limited. Approaches requiring the user to carry a tracking device [26] might not be suitable for long-term deployment either due to user inconvenience and privacy concerns.

Addressing the range limitations, mobile robot-based ULD approaches have been introduced [27], [28], [26] relying on cameras or distance sensors, still constrained by the robot's field of view. They localise users by mapping a room following semi-random patterns (exploration-only approach) [27], or incorporating the user's past frequented locations and "life rhythm" using Bayesian inference [29]. These approaches, although promising, are still limited to single room and single user localisation settings.

Formulating the localisation of non-stationary users in partially-known or unknown indoor environments as a path planning problem, the *"Where is Waldo?"* [16] style localisation models approach the question from the algorithmic side, focusing on modelling the long-term dynamics of object occurrences at given locations over time [19], aiming to increase the localisation's robustness. These works, however, treat this problem as a form of the Travelling Deliveryman Problem [30] focusing solely on the temporal aspect, without accounting for various localisation strategies' implication on user impressions and frustration [8], that cannot be defined by localisation time alone.

In attempt to account for the robot's effect on the user during interactions, user-aware interaction models often build on explicit user reactions (verbal or physical) to the robot's actions or implicit feedback signals. Such methods utilise social signal processing [31], analyse natural language responses [32] or other methods of modelling user affective states from implicit user behaviour expressed through multimodal data-streams (e.g. facial expression, body pose, gaze, gestures) [33], [34] or involuntary bodily functions (e.g. pulse or respiration) [35], [36] displayed by users during human-robot interaction scenarios. Prior works successfully inferred frustration, engagement or overall affective state (i.e. arousal-valence) of users [37], or used feedback signals to directly train learning models interpreting user impressions as a reward function [38] or evaluative feedback [39]. These methods are however, are constrained by the requirement for the robot to be in close proximity to the user to collect continuous (explicit or implicit) feedback reflecting their affective state during the interaction, and also complicated sensory requirements in some cases.

III. Frustration-Aware Assistive User Localisation

Problem Definition: The frustration-aware user assistive localisation problem consists of two components. Firstly, the user location estimation component: a given target user's location has to be estimated for a given discrete time interval (i.e. expected assistance interval). Secondly, the user localisation component: the agent has to navigate to the target user's true location, without inducing too much frustration in the target user (e.g. by spending longer than expected time with the localisation) while exploring the environment. To achieve this, the agent takes the location information of various users it encounters as an input, and produces a target user location estimate and a path to traverse as outputs, while accounting for the expected user frustration considering their individual characteristics, the task's importance and the proposed path.

Formally the frustration-aware user localisation problem can be defined as shown below:

Definition 1: A *user localisation problem* is a tuple $\langle U, T, G, P, I, \mathscr{F}, F(\cdot)\rangle$, where

- $U = \{u_1, u_2, ..., u_n\}$: set of human users.
- $T = [t_1, t_2, ..., t_n]$: set of discrete time slots.
- $G = (L, C)$ is a graph describing the agent's operational environment, consisting of location nodes $L = \{l_0, l_1, ..., l_n\}$ and connecting edges $C \subseteq \{\{l_1, l_2\} | l_1, l_2 \in$

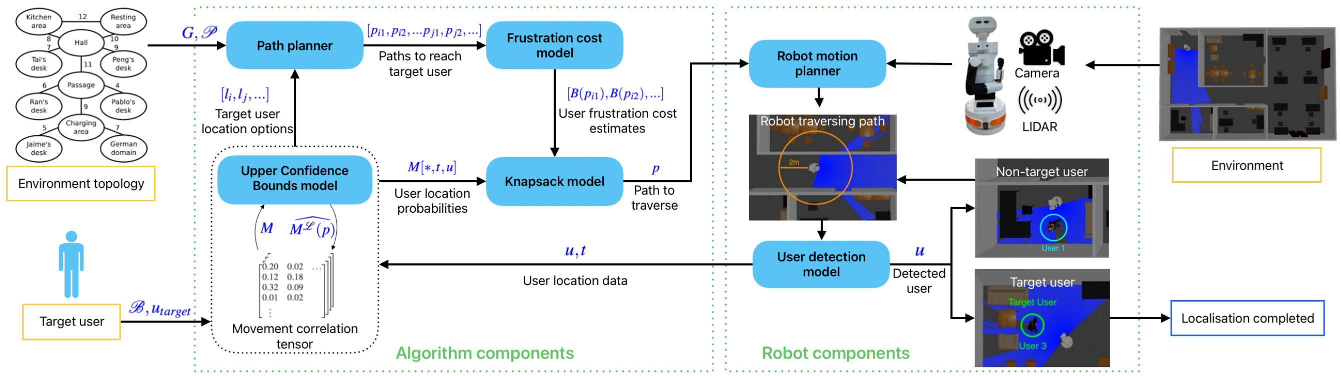

Fig. 1: Overview of the frustration aware user localisation framework's components. Upper Confidence Bounds model for updating learnt Movement Correlation Tensor values and estimating target locations, a psychological principles-inspired cost model for estimating user frustration from various robot actions, and knapsack optimisation methods for selecting paths.

$L, l_1 \neq l_2\}$ without "loops". Each edge $c \in C$ has a travel cost (weight).

- $P = \{\{l_0, l_i, l_j, \dots\} | \forall i, j \{l_i, l_j\} \in C\}$: a set of all possible paths or walks (any finite sequence of interconnected locations that can be traversed), starting from l_0.
- $I \in 1, 2, 3$: task importance (low, medium, high).
- $\mathscr{F} \in [1, 10]$: frustration tolerance level (low to high).
- $F(\cdot) : p, I, \mathscr{F} \to \mathbb{R}^+$ is the frustration cost function returning the estimated frustration cost of a movement.

Movement Correlation Tensor: The likelihood of a user to be located at a location at a time slot is represented as a 3-dimensional tensor $\mathscr{M}$, called the *movement correlation tensor*. Each row of the tensor corresponds to a location $l \in L$ and each column corresponds to a timeslot $t \in T$, with the third dimension representing the user $u \in U$. The values of $\mathscr{M}$ represent the true likelihood that user u is located at location l at timeslot t, with values $[0, 1]$, while M represents the learning model's estimation of $\mathscr{M}$. As they represent probabilities, the elements of each column of the tensor have a sum of 1. At each time slot t, the location of user u, denoted $l(u, t)$, is determined by random sampling with probabilities $\mathscr{M}[*, t, u]$.

User Localisation: The mobile robot localises users by traversing a chosen path $p \in P$ and exploring each location $l \in p$. Each movement between consecutive locations in p (i.e. from l_i to l_j where $l_i, l_j \in p$) has an associated

- frustration cost $F(\{l_i, l_j\}) > 0$.
- set of values $\{\mathbb{1}(l_j = l(u, t)) \, | u \in U\}$ at time slot t, where $\mathbb{1}$ denotes 1 if the predicate is true and 0 otherwise.

Put differently, after moving from l_i to l_j, the robot explores location l_j and observes whether any user $u \in U$ is at l_j. Then it updates the movement correlation tensor M based on its observations (i.e. $M[l_j, t, u] = \mathscr{L}(\mathbb{1}(l_j = l(u, t)))$ with $\mathscr{L}$ being an update function). We assume that once the robot is at location l_j, it can observe if any user is at l_j at the same time.

Frustration Cost Model: We designed our framework to be generic and compatible with different frustration cost models describing how the robot performing actions while traversing a path $p \in \mathscr{P}$ might affect users' frustration levels. For the purpose of running experiments, we utilised a frustration model focusing on the movement aspect of the task. It should be emphasised that from human psychological perspectives [40], [41], frustration or annoyance refers to the emotional state that arises when individuals encounter barriers (also known as frustration block) that prevent them from attaining a goal or satisfying their needs. Taking into the context of human-robot interaction (HRI), we can define an incident where the user is awaiting the robot's assistance (initialised with the localisation step) without the expected timely result as a cause for frustration. In other words, the frustration block refers to the delay in the robot's actions to respond to users in an expected time window. Previous psychological study [42] indicated that if the frustration block is not resolved, the frustration and resulting urge to aggression will become more dominant as the thwarting continues (Frustration-Aggression Hypothesis). In the early Human-Computer Interaction (HCI) work [43], the authors evidenced that the severity of frustration blocks, measured as the time lost and time to fix the problem, is correlated with the user's frustration level.

Building on the established psychological findings of the aforementioned works and interruption-based [44] frustration models, we define user frustration as

$$F = 1 - \alpha^{max(c(p) - t_{end}, 0)(0.25I^2 - 0.25I + 0.5)} / (1 - \alpha) \qquad (1)$$

where $\alpha = 1.21 - 0.02\mathscr{F}$ scaling ensures that the frustration function follows more linear or exponential pattern for more or less tolerant (i.e. $\mathscr{F}$) users respectively. The first exponential component models the robot's delay relative to the expected assistance interval's end (t_{end}), with $c(p)$ being the travel cost of p (i.e. duration of travel between the end points of path p in the given trial, quantified in seconds). The second exponential component quantifies the task importance (I) based scaling parameter. If the user's frustration level reaches its limit – we use an arbitrary limit value of $\mathscr{B} = 100$ representing 100% – the user ceases waiting, stops expecting the robot's help and the localisation task has failed.

Objective Function: Our goal is to localise a selected target user $u \in U$ at each time slot $t \in T$ by traversing a chosen path p, without exceeding the frustration limit $\mathscr{B}$ given the characteristics of the task and the user. To do this, we simultaneously need to estimate the true (oracle) movement

Fig. 2: Left: Topological map of Aruba environment [16]. Center: Simulation of Aruba environment by [16]. Right: Realistic office environment simulated in Gazebo including TIAGo and three users. Significant locations are marked, the blue shaded area shows the robot's rangefinder's visibility.

correlation tensor, denoted $\mathscr{M}$, by exploring a chosen path p. Then, we can update our estimate M with observations made using update function $\mathscr{L}$. Let the updated tensor be denoted $M^{\mathscr{L}}(p)$. Thus, we need to select an optimal path p^*:

$$p^* = \arg\min_{p \in P} \frac{D_{KL}(\mathscr{M} \mid M^{\mathscr{L}}(p))}{|U| \times |T| \times |L|} \; s.t. \; F(p, I, \mathscr{F}) \leq \mathscr{B} \quad (2)$$

Solution Concept: In our approach, we aim to approximate the oracle movement correlation tensor $\mathscr{M}$. Since $\mathscr{M}$ is not known in advance, the optimisation problem cannot be solved optimally. Our framework follows the same general structure for all variants: first, we initialise the movement correlation tensor M with random values. Then, at each consecutive time slot, we select a path p to traverse using a chosen heuristic, start moving along the path while observing encountered users at each location, and update our model based on the observation. Finally, once the target user is found, the robot can perform additional assistive tasks.

A. Upper Confidence Bounds model

The Upper Confidence Bounds (UCB) algorithm [45] performs continuous exploration and exploitation based on a set of confidence indexes, rather than in a stochastic manner. It is designed following the principle of "optimism in face of uncertainty", referring to the fact that based on the previously observed rewards, we can assign an optimistic upper confidence bound value to each potential action. This value will be an overestimation of the unknown true mean with a large probability, and with the number of actions taken, its value converges towards the true mean. These upper confidence bound values, or confidence indexes are defined as the sum of an exploitation and an exploration component. The former characterises how to best utilise the information that the model has collected so far, and is based on the reward history associated with the given action (i.e. mean reward so far, $\bar{x_j}$). The latter characterises the potential in exploring locations with little prior information, based on the current index value's uncertainty given by the model's exploration history and the variance of its past locations. This approach also enables the robot to explore locations with little past information. The model estimates the target user's location by choosing the location with the largest upper confidence index in the movement correlation tensor.

$$l = \arg\max_{l_j \in L} \; \bar{x_j} + \sqrt{\frac{\ln n}{n_j} \min\{1/4, V_j(n_j)\}} \quad (3)$$

where n_j is the number of attempts of j and V_j the variance.

Once a target user location estimate has been produced, we use a *knapsack model* [46] to identify the sequence of locations that maximise the probability of the target user is there. Here, each location is represented by an item within the knapsack. This item's value is the target user's probability of being there, and its weight is the frustration generated while travelling there from the previous location, with the frustration limit $\mathscr{B}$ being the knapsack's capacity. We used the following three heuristics to implement the knapsack based solutions:

- *One shot*: we solve the knapsack problem once to identify the path at once (i.e. all locations in the path that maximise the sum of expected probabilities of the target being there, without exceeding the frustration limit).
- *Multishot*: the knapsack problem can be solved multiple times generating the path sequentially. Each sub-path provides the path to the next most likely target location.
- *Shortest path*: it selects the shortest path to the estimated target position.

IV. Experiment Setup

To evaluate the proposed models, we designed two testing environments: (1) a realistic simulator used to mimic a typical multi-room office environment; and (2) an abstract simulator denoting the key navigation landmarks in an office space as a graph of partially interconnected nodes. Using the Gazebo office environment simulator, we simulate a TIAGo Steel robot interacting with different human users following behaviour patterns from established human behaviour datasets [18], [19]. In addition, the graph-based simulator allows us to test the proposed approach with a large number of simplified environments without environmental variables.

A. Office Environment Simulator

The simulated environment aims to mimic a realistic office workspace with private offices, open spaces, and communal areas (as shown in Figure 2) consisting of seven rooms (or locations). Each of these is populated with typical office

Task Characteristics		
Assistance Task	**Expected Assistance**	**Task Importance**
Transport parcel from reception	10:00-11:00	Low
Bring medicine and water	12:00-12:15	High
Deliver sandwich	12:30-13:00	Medium

User Characteristics		Model Status	
Target User	**Frustration Tolerance**	**Target Location Confidence**	**Environment Exploration Rate**
Alice	10	40%	30%
Bob	6	50%	70%
Charlie	8	80%	60%

TABLE I: Example assistive localisation task details.

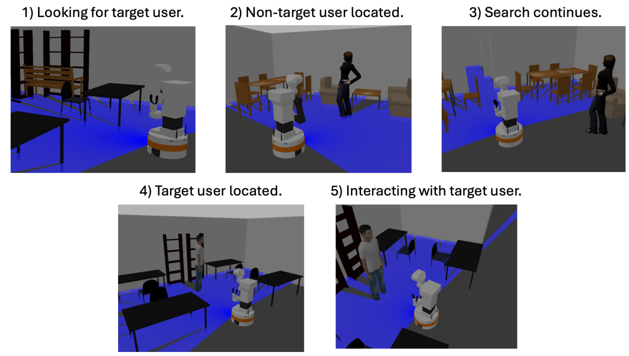

Fig. 3: Typical sequence of events when localising users in the Gazebo realistic office environment.

equipment such as desks, chairs, cupboards, coffee tables, and sofas in the communal area. This setup enables the formulation of correlations in individual users' movement patterns (e.g. jointly frequenting the breakroom around lunchtime) and potential accidental encounters. The various pieces of furniture ensure that navigation remains realistic since travel speeds and directions often have to be adjusted to accommodate to the environment (e.g. tight doorways or halls forcing the robot to slow down). As part of executing the localisation and assistance tasks, using the LIDAR and ultrasound sensors, TIAGo has to autonomously navigate the office space using a map of the various rooms' layout and the position of static pieces of furniture while being aware of dynamically moving users. Since our office environment contains multiple users (2-10 in our experiments), TIAGo has to identify them when localising. For this task, we assume that TIAGo can detect their identity once the user and the robot are in the same room. In our experiment, this setup is achieved symbolically: if TIAGo navigates to a location where its distance to a user is below a threshold, the simulator notifies TIAGo. This setup can be realistically fulfilled in real world deployments, since prior works show that users can be detected and identified with high confidence from this distance via facial recognition [47] or voice recognition [48].

B. Abstract Environment Simulator

Our abstract graph-based simulator aims to simulate the topology of typical office spaces. As illustrated in Figure 2, nodes represent significant locations, rooms in the environment, and weighted edges represent connections with distance information between neighbouring locations. The robot and users are located at the nodes (locations) during time slots and can move between nodes following the edges when necessary. The time of traversing the same edge may differ between experiments, as in real environments. This setting enables us to simulate human users' and the robot's movement patterns in a lightweight, customisable environment while still simulating TIAGo's and users' realistic motion characteristics. This approach allows us to rapidly generate a large number of environments with various sizes, topologies and interconnectivity rules, with magnitudes faster execution times enabling large scale experiments such as scalability analysis tests possible[1].

[1]The large scale experiments were performed using the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton.

C. Experimental Setup

Localisation Process: To localise the target user, TIAGo provides an estimated location of the target user using our UCB model and then generates a corresponding path to take using a chosen heuristic (oneshot, multishot, or shortest path). This path is then traversed by the robot while encountered users' locations are updated in the model's movement correlation tensor M. This information is used when performing localisation in the consecutive days as well.

User Behaviour and Characteristics: In their movements throughout the day, users follow behaviour patterns extracted from two established datasets: 1) "Aruba" dataset [18] originally collected by the Centre of Advanced Studies in Adaptive Systems (CASAS) as part of the WSU CASAS smart home project, further processed extracting user behaviour patterns as part of the STRANDS project [19]. 2) "Brayford" dataset [19] created by the Lincoln Centre for Autonomous System (LCAS) as part of the STRANDS project [19]. The two datasets contain user activities observed in household and office environments respectively over multiple weeks, making them ideal for modelling typical user behaviour patterns [16] suitable for simulations. The resulting room occupancy and user movement distribution data is used to populate each user's movement likelihood matrix formulating M, describing their likelihood of being at locations at certain times. At the beginning of each day, users' movements for the day were sampled from these matrices. All users are randomly assigned a frustration tolerance level $\mathscr{F} \in [1, 10]$ with 1 being the least tolerant and 10 being the most tolerant. This level remains constant throughout the experiments and is known to TIAGo in advance.

User Localisation Experiments: We observed scenarios where the robot assisted over 20 consecutive days, executing 8 assistive tasks at various times every day. User assistance tasks are assigned to the robot following the typical occurrence time of such tasks in Aruba or Brayford (e.g. deliver lunch request around noon), along with a task importance level (low, medium or high) and an expected assistance time interval ranging form 15 to 60 minutes (e.g. 12:00 to 12:30). TIAGo begins the localisation process from a fixed location (l_0) for each assistance request, where it has returned to at the end of the previous task to charge on standby. If the robot's behaviour results in the user's frustration level exceeding

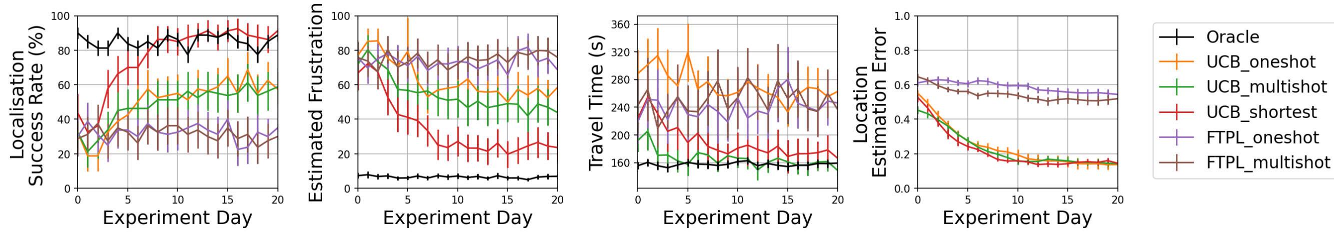

Fig. 4: Comparison of UCB and various benchmarks (FTPL, Oracle) in the realistic office environment, following "Aruba" user behaviour data. The plots display: 1) the model localisation & task execution success rates, i.e. the percentage of users successfully localised 2) estimated user frustration during assistance attempts 3) travel time until successful localisation or task failure 4) localisation error following the objective function. The displayed values are the mean of 20 experiments replicates, the error bars represent the confidence interval (1.96*SE).

their tolerance limit (e.g. 100%) considering factors such as their individual tolerance level or the task's importance, the user gives up on the robot's help and ceases to wait for the robot's assistance – resulting in a failed assistance attempt.

Benchmarks: The UCB models' one shot (*UCB_oneshot*), multishot (*UCB_multishot*) and shortest path heuristics (*UCB_shortest*) were evaluated along two benchmarks: A random baseline (*Random*), which randomly explores locations in the environment until the user's frustration allows; and two Follow-the-Perturbed-Leader (FTPL) [49] location estimation models (*FTPL_oneshot* and *FTPL_multishot*) based on multiarmed bandit theory's Follow-the-Leader algorithm [50], combined with one shot and multishot heuristics. The latter model is commonly used for structured online learning problems, as its implementation is comparable to solving the offline optimisation problem [51]. In addition, the results are compared against an oracle value, defined by a hypothetical scenario when the localisation agent's decisions are based on the oracle movement correlation tensor $\mathscr{M}$. All reported values are the result of 20 independent replicates.

Evaluation metrics: The error values for each estimation model's accuracy are calculated using the objective function as illustrated in Eq. 2. The success of the localisation process is characterised by the ratio of successful localisations over all attempts (%), the estimated user frustration level and the robot's travel time (s) till localisation or failure.

V. Results and Discussion

We evaluated the proposed models using both environments described in IV-C, using the "Aruba" [18], [19] and "Brayford" [19] datasets to model realistic user behaviours. Evaluations on both datasets yielded similar results, for the sake of presentation and readability, we present our results from the "Aruba" dataset experiments. First, we considered the models' user location estimation accuracy. Next, we evaluated the performance of the localisation process, following both experiment types, with additional focus on analysing TIAGo's navigation capabilities. Finally, we performed scalability analysis on the models' performance.

Location Estimation Error: The results presented in Figure 4.4 show the location estimation error of various models against the current day throughout the experiments. The solutions continuously improve before reaching a plateu, as

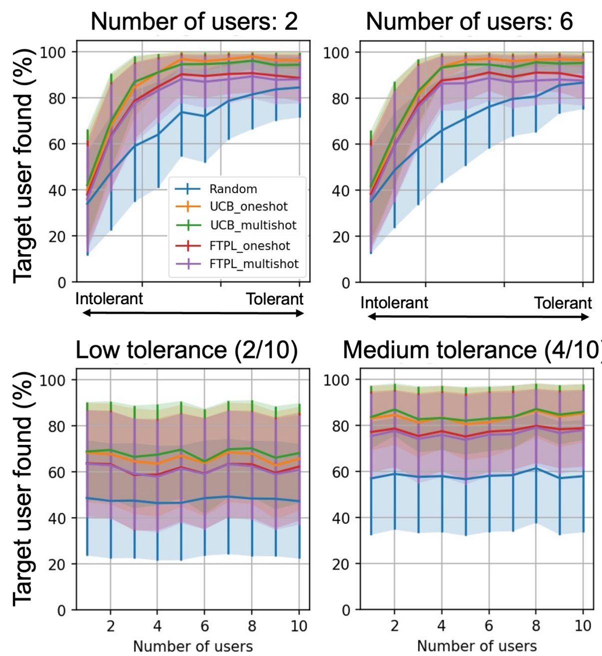

Fig. 5: Experiments conducting scalability analysis on the frustration tolerance (plots 1 and 2) and number of users (plots 3 and 4). The plots show the mean of 20 experiment replicates, the error bars represent the variance in target users found. The experiments were conducted in graph-based test suites, containing 3, 5 or 7 rooms environments each time.

illustrated by the decreasing error values as the days go on, while UCB models continuously outperform FTPL models.

Localisation Performance: Considering the overall localisation performance, Figure 4.1 shows that the UCB shortest localisation success rates are converging towards the oracle values with time, from around 40% success rate to the Oracle rate at around 90%, while the FTPL models' success rates only increase mildly under the given problem conditions. The UCB shortest path models perform the best, followed by UCB one shot and UCB multishot since the first model prioritises localising the target user more, while UCB one shot focuses on environment exploration and UCB multishot strives for a balance. The same tendencies show regarding estimated user frustration (Figure 4.2) with UCB models yielding up to 75% less frustration than FTPL models considering the mean difference over the experiment days, with UCB oneshot yielding approximately 20% less frustration. UCB multishot and UCB oneshot perform similarly. Regarding travel time till localisation or failure (see Figure 4.3), the UCB oneshot is on par with or even underperforms FTPL models in the first days, despite the former yielding less frustration, highlighting the importance of non-temporal factors in user-aware behaviour models. UCB oneshot is able to achieve frustration levels similar to UCB multishot, while allowing time for additional exploration with travel times around FTPL models. UCB multishot and shortest converge towards the Oracle travel times. The users were localised most days between an average of 160-200 seconds using the UCB shortest and multishot models (in comparison to the 1020 seconds necessary for a typical full traversal, or the mean 310 seconds necessary for the random baseline).

TIAGo's Performance: Regarding TIAGo's performance, it executed the navigation tasks accurately in most cases and adjusted its movement patterns to the environment (e.g. speed and direction adjustments near doorways or crowded environments). This also resulted in the different traversal paths chosen by various models yielding different travel time between the same start and end locations. However, in 10-15% of all cases, it did not navigate in the most optimal manner. Particularly in private offices or near doorways, the robot had a tendency of getting stuck or lost and took significant extra time (up to 30 seconds) to recover. Consequently, additional self-localisation steps were added reducing the duration of such events to under 5 seconds..

Scalability Analysis: Our scalability analysis (see Figure 5) in the abstract simulator has shown good scalability properties: with the increase of the user tolerance, the model performance (successful localisation rate) also increased until a plateau is reached near the oracle values, regardless of the environment's topology. The results also remain consistent as the number of participating users increases, with less than 5% difference in the mean target user found metric, illustrating that the models scale well. In both cases, the UCB models outperform both FTPL models and the random baseline.

VI. Conclusions and Future Work

In this work, we proposed a mobile robot based framework for localising and assisting users with tasks of various importance in an office environment behaving in a frustration-aware manner. The user frustration factor was taken into consideration by applying multiarmed bandit approaches and knapsack algorithms in addition to providing a psychology-based model of frustration suited to HRI settings without continuous user contact. We evaluated our approach following realistic user behaviour datasets in an office simulation using the TIAGo robot and a custom abstract simulation framework, with a particular focus on the framework's performance and scalability. Our results indicate that our proposed framework provides a suitable foundation for planning user localisation and assistance tasks in a frustration-aware manner in HRI environments without continuous user contact: we demonstrated localisation success rates and travel times converging towards oracle value (outperforming FTPL and random benchmarks) while continuously maintaining below tolerance limit user frustration; with our results indicating good scalability characteristics.

In future works, we aim to conduct in person HRI user studies evaluating both objective and subjective user impressions in response to interacting with frustration-aware robotic behaviour models, including settings without continuous user contact, serving as further verification of our current findings. Further on, attempting to enable a more realistic portrayal of user reactions, additional factors influencing user impressions or frustration in OAR contexts, different ways of modelling user frustration and deployment possibilities in other interaction scenarios are to be explored as well.

References

[1] F. Vicentini, "Collaborative Robotics: A Survey," *Journal of Mechanical Design*, vol. 143, no. 040802, 2020.

[2] D. Kragic, J. Gustafson, H. Karaoguz, P. Jensfelt, and R. Krug, "Interactive, Collaborative Robots: Challenges and Opportunities," in *IJCAI*, pp. 18–25.

[3] B. Gucsi, D. S. Tarapore, W. Yeoh, C. Amato, and L. Tran-Thanh, "To Ask or Not to Ask: A User Annoyance Aware Preference Elicitation Framework for Social Robots," in *IEEE/RSJ IROS*, 2020, pp. 7935–7940.

[4] A. Ghadirzadeh, X. Chen, W. Yin, Z. Yi, M. Bjorkman, and D. Kragic, "Human-Centered Collaborative Robots With Deep Reinforcement Learning," *IEEE RA-L*, vol. 6, no. 2, pp. 566–571, 2021.

[5] G. Nejat, Y. Sun, and M. Nies, "Assistive Robots in Health Care Settings," *Home Health Care Management & Practice*, vol. 21, no. 3, pp. 177–187, 2009.

[6] 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.

[7] 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.

[8] 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.

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

[10] K. Candon, Z. Hsu, Y. Kim, J. Chen, N. Tsoi, and M. Vázquez, "Perceptions of the helpfulness of unexpected agent assistance," in *ACM Human Agent Interaction*, 2022, pp. 41–50.

[11] C. Zhang, J. Chen, J. Li, Y. Peng, and Z. Mao, "Large language models for human–robot interaction: A review," *Biomimetic Intelligence and Robotics*, vol. 3, no. 4, p. 100131, 2023.

[12] S. I. A. P. Diddeniya, A. M. S. B. Adikari, H. N. Gunasinghe, P. R. De Silva, N. C. Ganegoda, and W. K. I. L. Wanniarachchi, "Vision based office assistant robot system for indoor office environment," in *International Conference on Information Technology Research*, 2018, pp. 1–6.

[13] M. Gouko and Y. Arakawa, "A coaster robot that encourages office workers to drink water," in *ACM Human Agent Interaction*, 2017, pp. 447–449.

[14] C. S. Franklin, E. G. Dominguez, J. D. Fryman, and M. L. Lewandowski, "Collaborative robotics: New era of human–robot cooperation in the workplace," *Journal of Safety Research*, vol. 74, pp. 153–160, 2020.

[15] W. Wang, N. Liu, R. Li, Y. Chen, and Y. Jia, "HuCoM: A Model for Human Comfort Estimation in Personalized Human-Robot Collaboration," in *ASME Dynamic Systems and Control Conference*, 2018.

[16] T. Krajník, M. Kulich, L. Mudrová, R. Ambrus, and T. Duckett, "Where's waldo at time t ? using spatio-temporal models for mobile robot search," in *IEEE ICRA*, 2015, pp. 2140–2146.

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

[18] 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.

[19] T. Krajník, J. P. Fentanes, O. M. Mozos, T. Duckett, J. Ekekrantz, and M. Hanheide, "Long-term topological localisation for service robots in dynamic environments using spectral maps," in *IEEE/RSJ ICRA*, 2014, pp. 4537–4542.

[20] E. Ahvar, N. Daneshgar-Moghaddam, A. M. Ortiz, G. M. Lee, and N. Crespi, "On analyzing user location discovery methods in smart homes: A taxonomy and survey," *Journal of Network and Computer Applications*, vol. 76, pp. 75–86, 2016.

[21] F. Furfari, A. Crivello, P. Barsocchi, F. Palumbo, and F. Potorti, "What is next for Indoor Localisation? Taxonomy, protocols, and patterns for advanced Location Based Services," in *IEEE International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, 2019.

[22] P. Ssekidde, O. Steven Eyobu, D. S. Han, and T. J. Oyana, "Augmented CWT Features for Deep Learning-Based Indoor Localization Using WiFi RSSI Data," *Applied Sciences*, vol. 11, no. 4, p. 1806, 2021.

[23] X. Liang and J. Deng, "Detection of stationary human target via contactless radar networks," *Journal of Ambient Intelligence and Humanized Computing*, vol. 10, no. 8, pp. 3193–3200, 2019.

[24] K. Ngamakeur, S. Yongchareon, J. Yu, and S. U. Rehman, "A Survey on Device-free Indoor Localization and Tracking in the Multi-resident Environment," *ACM Computing Surveys*, vol. 53, no. 4, pp. 1–29, 2021.

[25] D. Yang, W. Sheng, and R. Zeng, "Indoor human localization using PIR sensors and accessibility map," in *IEEE International Conference on Cyber Technology in Automation, Control, and Intelligent Systems (CYBER)*, 2015, pp. 577–581.

[26] J.-G. Lee, M.-S. Kim, T.-M. Hwang, and S.-J. Kang, "A mobile robot which can follow and lead human by detecting user location and behavior with wearable devices," in *IEEE International Conference on Consumer Electronics (ICCE)*, 2016, pp. 209–210.

[27] H. Kudo, T. Koizumi, T. Nakamura, M. Kanoh, and K. Yamada, "Behavior Model for Hearing-Dog Robot," in *IEEE Joint International Conference on Soft Computing and Intelligent Systems (SCIS) and International Symposium on Advanced Intelligent Systems (ISIS)*, 2016, pp. 260–265.

[28] H. Kudo, S. Tanaka, Y. Yoshida, T. Nakamura, M. Kanoh, K. Yamada, D. Oiwa, Y. Iwahori, and S. Fukui, "Estimation of user location for hearing-dog robot using past experience," in *Joint World Congress of International Fuzzy Systems Association and International Conference on Soft Computing and Intelligent Systems (IFSA-SCIS)*. IEEE, 2017, pp. 1–5.

[29] S. Furuta, T. Nakamura, Y. Iwahori, S. Fukui, M. Kanoh, and K. Yamada, "Consideration of Life Rhythm for Hearing-Dog Robots Searching for User," in *IEEE Technologies and Applications of Artificial Intelligence (TAAI)*, 2018, pp. 102–105.

[30] F. Afrati, S. Cosmadakis, C. H. Papadimitriou, G. Papageorgiou, and N. Papakostantinou, "The complexity of the travelling repairman problem," *RAIRO - Theoretical Informatics and Applications - Informatique Théorique et Applications*, vol. 20, no. 1, pp. 79–87, 1986.

[31] 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.

[32] 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 International Conference on Multimodal Interaction (ICMI)*, 2014, pp. 42–49.

[33] M. Stiber, R. Taylor, and C.-M. Huang, "Modeling human response to robot errors for timely error detection," in *IEEE/RSJ IROS*, 2022, pp. 676–683.

[34] 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 *ACM International Conference on Multimodal Interaction (ICMI)*, 2017, pp. 181–188.

[35] 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.

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

[37] 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.

[38] 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*, J. Kober, F. Ramos, and C. Tomlin, Eds., vol. 155, 2021, pp. 604–626.

[39] 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.

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

[41] J. S. Brown and I. E. Farber, "Emotions conceptualized as intervening variables—with suggestions toward a theory of frustration." *Psychological Bulletin*, vol. 48, no. 6, pp. 465–495, 1951.

[42] N. E. Miller, "I. The frustration-aggression hypothesis." *Psychological Review*, vol. 48, no. 4, pp. 337–342, 1941.

[43] 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.

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

[45] 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.

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

[47] M. Ye, J. Shen, G. Lin, T. Xiang, L. Shao, and S. C. H. Hoi, "Deep learning for person re-identification: A survey and outlook," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 6, pp. 2872–2893, 2022.

[48] S. Argentieri, P. Danès, and P. Souères, "A survey on sound source localization in robotics: From binaural to array processing methods," *Computer Speech & Language*, vol. 34, no. 1, pp. 87–112, 2015.

[49] T. Lattimore and C. Szepesvári, *Bandit Algorithms*, 1st ed. Cambridge University Press, 2020.

[50] J. Hannan, "Approximation to Bayes risk in repeated play," *Contributions to the Theory of Games*, vol. 3, pp. 97–139, 1957.

[51] B. Hao, H. Du, J. Zhao, J. Zhang, and Q. Wang, "A Path-Planning Approach Based on Potential and Dynamic Q-Learning for Mobile Robots in Unknown Environment," *Computational Intelligence and Neuroscience*, vol. 2022, pp. 1–12, 2022.