



A Social Robot Learning to Facilitate an Assistive Group-Based Activity from Non-expert Caregivers

Wing-Yue Geoffrey Louie^{1,2} · Goldie Nejat²

Accepted: 7 January 2020 / Published online: 18 January 2020
© Springer Nature B.V. 2020

Abstract

Socially assistive robots are a promising technology for supporting residential care facilities to provide stimulating recreational activities to residents in group settings. In order for caregivers to teach robots customized recreational activities for residents in their facilities, these robots need to be able to learn such activities from non-experts. In this work, we present a novel learning from demonstration system that allows socially assistive robots to learn customized group recreational activities from caregivers and facilitate these activities with users. We validate the usability and effectiveness of the proposed system by conducting a robot teaching study with caregivers and the Tangy robot at a local residential care facility. The caregivers found the learning system easy to use, experienced moderately low perceived workload, and were able to successfully teach Tangy the game of Bingo. Once Tangy learned the game, it autonomously facilitated Bingo games with elderly residents. The residents found the robot behaviors, personalized by the caregivers, both helpful and entertaining. Furthermore, they enjoyed playing Bingo with Tangy and would participate in future games.

Keywords Socially assistive robots · Learning from demonstration · Non-expert robot teachers · Group-based activities · Caregivers and older adults

1 Introduction

Socially assistive robots (SARs) are being integrated into residential care settings to assist in engaging older adults in group-based recreational activities [1, 2]. Such stimulating activities have the potential to maintain or improve older adult health by reducing the risk of cognitive and physical decline, heart conditions, and depression [3]. However, current robots are limited to the activities pre-programmed on the robots by roboticists. Non-expert robot teachers (i.e. caregivers) cannot customize activities based on user needs. This limits the efficacy of robots as older adults have a diverse set of recreational needs due to differences in physical and cognitive health, functional capabilities, and activity

preferences [4]. Caregivers frequently adapt activities from their traditional counterparts to account for the capabilities of older adults [4]. Such customization of activities can improve activity engagement and the overall moods of older adults [5].

Our research focuses on developing SARs that can autonomously facilitate recreational activities with older adults in residential care settings to provide both cognitive and social stimulation. Our current efforts are towards developing robots that are capable of learning customized group recreational activities from non-experts (i.e. caregivers) in residential care settings and implementing these activities with older adults. Such customization capabilities will allow caregivers to improve engagement of older adults in these activities by adapting them to the needs and preferences for their facilities.

In this paper, we present the development and implementation of a unique learning from demonstration (*LfD*) system for SARs to learn customized group recreational activities from non-expert demonstrations. Herein, we define non-expert teachers as individuals inexperienced with programming robots. The main contributions of this work are: (1) the development of a *LfD* system architecture capable of

✉ Wing-Yue Geoffrey Louie
louie@oakland.edu

¹ Intelligent Robotics Laboratory, Department of Electrical and Computer Engineering, Oakland University, Rochester, MI, USA

² Autonomous Systems and Biomechanics Laboratory, Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, ON, Canada

learning the structure of non-sequential multi-step social activities without requiring explicit activity rules from teachers and a model of demonstrators' non-deterministic activity facilitation behavior; (2) a robot teaching study with caregivers using the proposed *LfD* architecture to teach a social robot a cognitively stimulating recreational activity in order to investigate usability, workload and overall user experience; and (3) a human–robot interaction (HRI) study with older adults using the learned activity and robot behaviors to determine if the robot's personalized activity behaviors were helpful and engaging.

2 Robot Task Learning from Demonstration

LfD approaches can be categorized into: action-level learning—low-level manipulation tasks, and task-level learning—using low-level actions to perform high-level complex tasks [6]. Our research focuses on task-level learning, which is achieved by learning state-action policies by mapping world states to primitive robot actions. Namely, a world state refers to all the states (e.g. robot state, object state, person state) relevant to a given task. Compared to end-user programming approaches such as those presented in [7], *LfD* approaches do not require any understanding of programming concepts (e.g. loops, functions, conditional statements, etc.) by the non-expert teachers. Instead the challenge of designing the programming logic for a task is addressed solely by the learning system.

In general, task-level learning using *LfD* has mainly focused on teaching robots to perform *physical tasks* such as assembling a table [8] or rotating car tires [9]. These physical tasks can be accomplished sequentially as the world is often static, only behaviors of the robot modify the world state, and human demonstrated behaviors are often deterministic with one behavior always being executed in the specific state [10].

Learning social tasks from demonstration is a new emerging research area. It is unique as the world is usually dynamic due also to the ability of the behaviors of the users engaged in the interactions to alter the world state. In a social task, these user behaviors are often due to changes in the user's intent, affect and/or needs. Therefore, the robot must be able to adapt in real-time to these different behaviors. Furthermore, the behaviors demonstrated by caregivers are often non-deterministic when facilitating social activities. For example, the caregivers in residential care facilities spontaneously provide encouragements, jokes, or instructions during a social activity with older adults to promote social engagement or encourage participation in the stimulating activity [11]. Hence, a robot learning to facilitate social tasks will need to model such non-deterministic behavior.

To-date, a handful of work has focused on addressing some of the challenges associated with robots learning social tasks from human demonstrations. These include: (1) our previous work on using *LfD* to have a robot learn group recreational activities from demonstrations conducted in simulation by non-expert teachers [12], (2) an interface designed for non-experts to develop policies for a service robot in a smart home [13], and (3) a methodology for a robot to learn social affordances from interactions to accomplish a physical task in collaboration with a human partner [14]. In our previous work, the social robot Tangy learned the group recreational activity Bingo from teleoperation-based demonstrations performed in simulation by university students. A teacher demonstrated a task by controlling the robot's pre-determined behaviors during a Bingo game simulation. The prior system was not capable of learning the individual behavior motion trajectories or speech from human demonstrations and could not learn a model of the non-deterministic activity facilitation behaviors of a demonstrator. In the present work we have extended the system architecture to enable non-expert teachers to demonstrate new individual behaviors to the robot and learn non-deterministic activity facilitation behaviors from their demonstrations. We also focused on evaluating the system at a residential care facility with non-expert caregivers, who are the targeted user group for our *LfD* system. Namely, we evaluated their perceptions on the usability, user experience, and workload of the proposed learning from demonstration system. Furthermore, we evaluated the performance of the *LfD* system for autonomously facilitating an activity with older adult residents after it has learned the activity from caregiver demonstrations.

In [13], non-expert teachers developed task policies for the Care-O-bot 3 through a GUI to have the robot perform simple *single step* assistive tasks such as providing reminders to older adults in a smart home. Each teacher used the GUI to generate a task policy by providing explicit rules for the task. These inputs were then used to model a task using IF–THEN rules. The robot teaching system was evaluated with young and old teachers. Participants thought the system was both useful and usable. The approach presented in [13] requires teachers to provide explicit rules, which places the challenge of creating accurate task policies in the hands of the non-expert teachers. This can be an especially complex process for non-sequential multi-step tasks as teachers would need to have an expert understanding of how a task is modeled and learned by the robot [6]. Furthermore, the aforementioned system is deterministic in that teachers can only assign one behavior to each world state, and therefore the execution of multiple different behaviors in a single world state is not possible.

In [14], the iCub robot learned social affordances from human interactions to enable it to socially interact with a human to accomplish an object placement task. Namely, a

study was conducted where the iCub robot utilized social affordances, such as socially requesting for human assistance with a set of verbal behaviors (e.g. “pass me”, “push left”, “sit down”), and participants physically responded to the robot’s requests. During the interaction the robot observed the effects verbal requests had on objects (e.g. change in location) and the participant (e.g. change in body orientation). SVMs were then trained from the observations to predict the effects of verbal behaviors on objects or a human. The learned SVMs were then used to plan the set of robot object manipulation and verbal behaviors to accomplish an object placement task in collaboration with a human. However, this approach requires a distinct one-to-one relationship between a robot behavior and a world state, namely a deterministic relationship. This limits its applicability to our assistive application because behaviors demonstrated by caregivers are often non-deterministic when facilitating social activities. Furthermore, each behavior in [14] is assumed to only have a single possible state transition but during a group recreational activity the execution of a behavior can have multiple potential state transitions.

In this work, we present a novel approach for robots to learn a policy for a non-sequential multi-step social task, such as a group recreational activity, from a teacher’s demonstrations without requiring a teacher to provide explicit rules. Namely, a dataset of world state-behavior pairs is captured during a teacher’s demonstration to learn a classifier where the world states are the features being classified and the behaviors are the class labels. The observed dataset is used to infer the relevant states for an activity and the order in which individual states are evaluated to determine the behavior a robot should execute during an activity. Furthermore, the system is capable of learning a policy to handle non-deterministic demonstrations by a single teacher. Namely, our system learns a policy that models the probability at which each behavior is executed in that world state, allowing the system to handle non-deterministic interactions with users while autonomously facilitating an activity when there are multiple appropriate behaviors that can be used in any given world state.

3 Proposed Robot Learning from Demonstration System Architecture

The objective of our *LfD* system architecture is to allow non-expert teachers from residential care facilities to easily teach and customize a multi-step activity and its corresponding behaviors to SARs. Towards this aim, we have developed a system that uses both teleoperation-based *LfD* for learning the activity structure and external observation based *LfD* for learning the activity behaviors. Herein, a behavior refers to the specific speech and motion trajectories a robot

implements to achieve a desired effect during an activity. For example, a greeting behavior would include the speech “hello” and a waving hand motion with the desired effect of starting the activity with a group of users. An activity then refers to the high-level social task (e.g. Bingo, Trivia) which requires decisions on what behavior a robot needs to implement to accomplish the task. Namely, we model an activity as a set of rules which defines the appropriate behavior to implement in a given world state. These definitions are described in detail in Sect. 3.1.2.

The proposed system architecture consists of two sub-systems, Fig. 1: ① *the demonstration sub-system*, and ② *the interaction sub-system*. The demonstration sub-system allows a teacher to demonstrate the facilitation of an activity. A keyboard, mouse, and 3D sensor are used as ③ *demonstration inputs* from the teachers. The inputs are used with the ④ *demonstration GUI* to demonstrate to the robot the behaviors and structure of an activity. The demonstrated activity trajectory is then used by ⑤ *the activity learning module* to learn an activity policy which defines the mapping of world states to behaviors. Behavior demonstrations are used by ⑥ *the behavior learning module* to learn individual behavior policies which define the motion trajectory for the robot’s two arms and its speech for a behavior. The interaction sub-system then utilizes the learned activity policy and the learned behavior policies to autonomously facilitate an activity. Namely, ⑦ *sensory* information is used to perceive the world during an activity, which is then classified by ⑧ *the identification of world state parameters module* into distinct

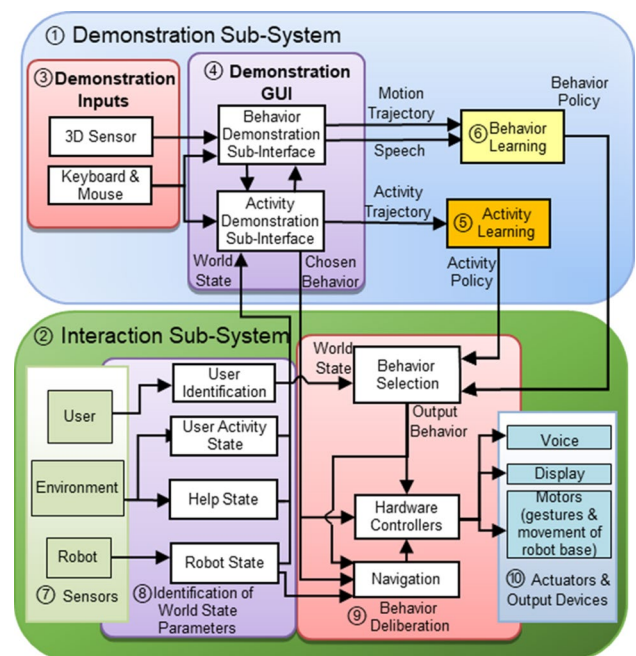


Fig. 1 *LfD* system architecture for robot activity learning

world states. The identification of world state parameters module identifies individual users, user activity states, help states, and the robot's state. Based on the world state, $\textcircled{2}$ the behavior deliberation module determines an appropriate robot behavior to execute, which is sent to the hardware controllers and navigation sub-modules to execute using the robot $\textcircled{10}$ actuators and output devices.

3.1 Demonstration Sub-system $\textcircled{1}$

The demonstration sub-system consists of three main modules:

3.1.1 Demonstration GUI $\textcircled{4}$

The demonstration GUI is utilized by a teacher to demonstrate an activity and consists of two components: (1) the behavior demonstration sub-interface, and (2) the activity demonstration sub-interface.

A teacher demonstrates the activity by setting up activity scenarios expected to occur during the facilitation of the activity. Herein, we refer to these activity scenarios as world states. The teacher then uses the activity demonstration sub-interface to choose the appropriate robot behavior to execute according to these scenarios. During an activity demonstration, the robot identifies the world states using its sensory information and the robot behaviors chosen by a teacher. The observed sequence of demonstrated world states and behaviors, referred to as the activity trajectory, is input into the activity learning module.

The behavior demonstration sub-interface is utilized by a teacher to demonstrate the gestures and speech for a behavior. Each behavior demonstration begins with a teacher classifying the behavior he/she is demonstrating. The teacher then inputs the desired speech for the behavior. Then he/she performs the desired gesture for the behavior, so that the robot can mimic him/her. The 3D sensor mounted on the robot is used to capture skeleton joint positions as the teacher performs a gesture. We refer to the sequence of observed teacher poses as the motion trajectory. The input speech and the motion trajectory demonstrated by the teacher are utilized in the behavior learning module.

3.1.2 Activity Learning $\textcircled{5}$

The activity learning module learns the world state to robot behavior mapping for an activity using the demonstrated activity trajectory.

World Model We model the world as a set, W , that consists of a specific instance, z , of the robot (R), user (U), and activity (A) states: $W^z = \{R^z, U^z, A^z\}$. We have previously determined that these states generalize across group activities and are significant for determining the facilitator's

behavior based on our focus group study in [15] as well as preliminary observations of various recreational activities (e.g. Trivia, Bingo) being facilitated by the caregivers.

Robot Model A robot, R , is a physical entity that can interact with the world using its behaviors. Each robot behavior, b , refers to the high-level goal of the robot and the behavior policy, π_b , which is a specification of the low-level actuation, speech, and location of the robot to accomplish this goal. For example, a robot behavior could be "Greeting a User" and the behavior policy specifies that the robot should wave its right arm and say "Hi" for this behavior. Herein, we define the robot's behavior policy as the speech the robot says (sph), the joint angles for the robot's two arms (θ_m), and the 2D location in the environment it should navigate to (l_b): $\pi_b^i = \{sph^i, \theta_m^i, l_b^i\}$, where i is a particular robot behavior. The group of learned and default robot behaviors, B , define all the possible behaviors a robot can implement: $B = \{b_1, b_2, \dots, b_n\}$ where n is the total number of robot behaviors. The state of the robot is defined by the robot's current 2D location (l_r) and the current behavior it is executing (b): $R^q = \{b^q, l_r^q\}$, where q is an instance of the robot state.

User Models Each user participating in the group recreational activity has his/her own name (ID), individual user activity state (s_{ua}), help state (s_h), and 2D location within the world (l_u): $u^j = \{ID^j, s_{ua}^j, s_h^j, l_u^j\}$ where j is a particular user. Herein, the help state is a binary state which defines whether a user is or is not requesting for assistance during the activity. The user activity state refers to activity specific conditions for each individual user that a facilitator needs to monitor during an activity. For Bingo, the user activity state is defined by the state of a user's Bingo card: Occluded, Bingo, Incorrectly Marked, or Missing Numbers. A total of m users can participate in an activity: $U = \{u^1, u^2, \dots, u^m\}$.

Activity Model An activity, A , is defined by the multiple distinct states that it can be in: $A = \{a^1, a^2, \dots, a^o\}$ where o is the total number of activity states for a particular activity. Namely, an activity state can be defined as a function of the users' states (U), robot state (R), and time step k : $a^c = f(k^c, U^c, R^c)$, where c is a particular activity state.

Activity Trajectory The activity trajectory is the sequence of observed world state-behavior pairs during a teacher's demonstration of the activity. Namely, an activity trajectory is modelled as: $T = \{b_1, W_1\} \rightarrow \{b_2, W_2\} \rightarrow \dots \{b_p, W_p\}$, where p is the total number of world states and behaviors observed. The activity trajectory is used as input into our activity learning algorithm to train a supervised learning classifier where each individual world state-behavior pair, $\{b_k, W_k\}$, in the observed trajectory is utilized as a training sample.

Activity Learning Algorithm In order to learn the activity, we take a supervised learning approach to train a classifier with the activity trajectory T demonstrated by the teachers.

Since we expect the demonstrations by the non-expert teachers will be non-deterministic, we utilized a classifier which can provide posterior probabilities for multiple robot behaviors for a given input world state. Namely, the input to the Random Forest classifier is the current world state instance (W^c) which includes the current robot state (R^c), user state (U^c), and activity state (A^c). The final output of our classifier provides posterior probabilities for each of the possible behaviors (B) that the robot can implement in the next world state. In this system, we utilize a random forest classifier because it has been shown to be the most accurate general-purpose classifier among the aforementioned classifiers and can learn with minimal training data in high-dimensional feature spaces [16–18]. These properties of random forests will enable our system architecture to be scalable to different group recreational activities with potentially larger state spaces and require minimal training data from the demonstrator. Namely, individual decision trees are first learned by sampling a set of world state-behavior pairs from the activity trajectory and using a binary recursive partitioning procedure to determine decision rules that best predict robot behaviors based on world states for the sampled data. We use the Gini Index metric for evaluation of a rule:

$$\text{Gini Index} = 1 - \sum_n \text{prob}_{b_i}^2, \quad 0 \leq \text{Gini Index} \leq 1, \quad (1)$$

where candidate rules with Gini Index values closer to 0 provide better classification of the behaviors to the world states observed in the training data. Decision trees are then generated from these candidate rules. These learned decision trees are utilized to facilitate an activity and adapt the robot's behaviors to the users in real-time. Namely, during the activity, the current world state, W^c , is identified via the robot's sensors and used as input into the decision tree in order to classify the appropriate behavior b^i for the robot to execute.

A random forest classifier for an activity consists of multiple learned decisions trees, which are used to vote on input world states to determine the probability distribution (Ω_B) over the set of behaviors B . Learning the probability distribution of behaviors in a given world state allows the robot to model when a teacher demonstrates multiple different behaviors in a single world state. During an activity the robot can then use the learned policy to probabilistically choose amongst the possible behaviors in a world state. Herein, the final learned random forest refers to the activity policy (π_a): $\pi_a(W^c) = \Omega_B$.

3.1.3 Behavior Learning ⑥

The behavior learning module allows teachers to teach the robot new behaviors as well as for them to modify default robot behaviors. We use external observation based *LfD* as

it provides a natural mode of HRI for non-experts to teach robot arm gestures [19]. Speech input is provided via a keyboard, which is also a familiar mode of interaction for teachers.

The input speech and motion trajectory are used by the behavior learning module to define the behavior policy (π_b) for a behavior. Namely, the speech input provided by a teacher using the behavior demonstration sub-interface is used to define the behavior speech (*sph*). For the joint angles for the robot's two arms (θ_m), we utilize a geometric-based inverse kinematics approach we developed in [20] to obtain the mapping of teacher joint angles to robot angles based on the teacher's demonstrated motion trajectory. We use a bounding-volume based approach to check for robot self-collisions when the joint angles are mapped to the robot due to differences in demonstrator and robot embodiment. Given a self-collision is detected, we optimize for a set of collision-free joint angles with the minimum difference to the demonstrated joint angles. Learned behavior policies are then used by the behavior selection module to define the appropriate speech and robot joint positions to execute when a behavior is implemented.

3.2 Interaction Sub-system ②

The interaction sub-system is used to determine the world state using the user, environment, and robot sensors, and physically implement the learned activity behaviors on the robot according to the activity and behavior policies.

3.2.1 Identification of World State Parameters ⑧

User identities (*ID*) are determined utilizing the user sensors. User help states (s_h), user activity states (s_{ua}), and user 2D locations (l_u) are identified using the environment sensors. A robot's 2D location (l_r) in a room is identified using robot sensors. These world state parameters are used to define the world state instance W^c and used as inputs to the behavior selection sub-module in the behavior deliberation module.

3.2.2 Behavior Deliberation ⑨

The behavior deliberation module determines the appropriate behavior to execute during an activity and implements behaviors by sending the low-level commands to the robot's actuators and output devices. Namely, it uses the world state instance W^c as an input in the learned activity policy function $\pi_a(W^c)$ to identify the mapped probability distribution (Ω_B) over the set of behaviors. We then sample over the mapped probability distribution to identify the behavior to implement. The robot behavior is then implemented using the navigation sub-module to plan

the robot's path and hardware controller sub-modules to implement the speech, visual information, arm trajectories, and wheel velocities.

4 Learning to Facilitate a Group Activity

In this work, our objective is to have non-expert caregivers from residential care facilities teach and customize a multi-step group recreational activity to the SAR, Tangy. As an example of such an activity, we selected the game of Bingo as it is a popular activity played at care facilities, which is often customized by staff to meet the needs as well as capabilities of residents [21]. These customizations are typically utilized to support residents with different cognitive capabilities and/or increase residents' engagement in the activity. For example, healthcare professionals working with a group of older adults with dementia often provide reminders on the rules of the game throughout the activity and occasionally check players' cards to identify mistakes or potential winning conditions [21]. Similarly, social utterances (e.g. jokes or facts or encouragement) focusing on the preferences of the residents are also utilized frequently throughout an activity to create a social atmosphere between older adults in the group and increase their engagement. Learning the scenarios these behaviors are executed in and the probability at which they are executed during an activity is important because if these behaviors are too frequent participants may become frustrated or disengaged but if behaviors such as reminders aren't provided frequently enough some individuals may not be capable of participating in the activity. Hence, the objective of our system is to learn both the structure of the activity and the non-deterministic activity facilitation behaviors that are unique to each facility. In [22], it was noted that at least 24 variations of Bingo were played among only four different facilities. Furthermore, Bingo has been shown to provide a number of benefits to older adults such as improved memory, recall, and recognition functions, and opportunities for social engagement with other players [23, 24].

4.1 Bingo

A Bingo game consists of Tangy standing at the front of a room and randomly calling out numbers from 1 to 75 while players mark these numbers on their cards with red markers. Each player sits behind a table with a Bingo card and an assistance request device, Fig. 2. During the game, a player can press the green button on the assistance request device to ask the robot for help.




Fig. 2 Tangy and the Bingo scenario

4.2 The Socially Assistive Robot Tangy

To facilitate recreational activities, the Tangy robot, Fig. 2, uses multi-modal interactions. Tangy is a human-like social robot that can mimic human gestures using three degrees-of-freedom (DOFs) for each shoulder and one DOF for each elbow. The robot also has a two DOF neck. Tangy is able to verbally interact with users using its synthesized voice. It also uses its chest mounted tablet to display both images and text. Tangy retrieves world state parameters including activity, user, and environment information using multiple sensors. The robot's sensors include: a 2D Logitech C920 camera mounted on top of its head, an ASUS Xtion IR sensor mounted behind the robot, a 2D Axis M1031-W camera in the robot's right eye, a URG-04LX-UG01 laser range finder on the robot's base. Furthermore, an additional ASUS Xtion IR sensor mounted on the robot's chest is used to identify teacher poses during behavior demonstrations.

4.3 Bingo Learning Scenario

The overall goal is to have a teacher demonstrate to Tangy the structure of the Bingo activity. A teacher can either create new behaviors, use default behaviors or customize the default behaviors. These behaviors can be used to teach the structure of the game to the robot. Examples of default Bingo behaviors are shown in Table 1. To create a new behavior, a behavior name is first given. Then the robot's gestures, speech, and 2D location can be customized by the teacher. This customization can also be done for the default behaviors.

During a Bingo activity demonstration, a teacher utilizes the activity demonstration sub-interface to teach the robot the structure of the game using the behaviors.

Table 1 List of example robot bingo behaviors

Robot behavior	Default behavior
Greet players	“Hi! My name is Tangy. I am so excited to play Bingo with you today” (Robot waves at players)
Tell a joke	“Why wouldn’t the shrimp share his treasure? Because he was a little shellfish” (Robot brings its hand to its mouth and giggles)
Call out a Bingo number	“The next number is: B-12” (Robot points to number displayed on its tablet)
Encourage a player	“Wow, you are close to getting Bingo!”
Ask a player to remove incorrectly marked numbers on Bingo card	“Oops! You have some misplaced markers. Please remove these markers from the following numbers on your card: B-5, O-65” (Robot points to numbers displayed on its tablet)
Ask a player to mark missing numbers on Bingo card	“I think you have some missing markers. Can you place markers on the following spots? B-1” (Robot displays numbers on its tablet)
Celebrate a winning Bingo card	“Congratulations, you won Bingo!” (Robot raises both its arms straight up and sways them side to side in a celebration dance)
Ask a player to show Bingo card	“Please move your Bingo card forward”
Navigate to a player who requires help	(Robot physically navigates to a player that requires help)
Navigate to the front of the room	(Robot physically navigates back to the front of the room)
Say goodbye to players	“Well that’s it! Wasn’t that fun? Thank you very much for playing Bingo with me. I had so much fun! Goodbye!” (Robot waves at players)

4.3.1 Activity Demonstration Sub-interface

The activity demonstration sub-interface, Fig. 3, has three main features: (1) robot command center to demonstrate an activity to Tangy by having the teacher teleoperate Tangy and control the robot’s behaviors through a complete game (red box); (2) sensor views which includes the robot view, help view, and card view (green box); and (3) information section which provides the list of activity scenarios

that have already occurred during the demonstration and the robot behavior that was implemented in each scenario (blue box).

4.3.2 Behavior Demonstration Sub-Interface

The teacher uses the behavior demonstration sub-interface, Fig. 4, to customize or create behaviors. We designed our behavior demonstration sub-interface to guide the teacher

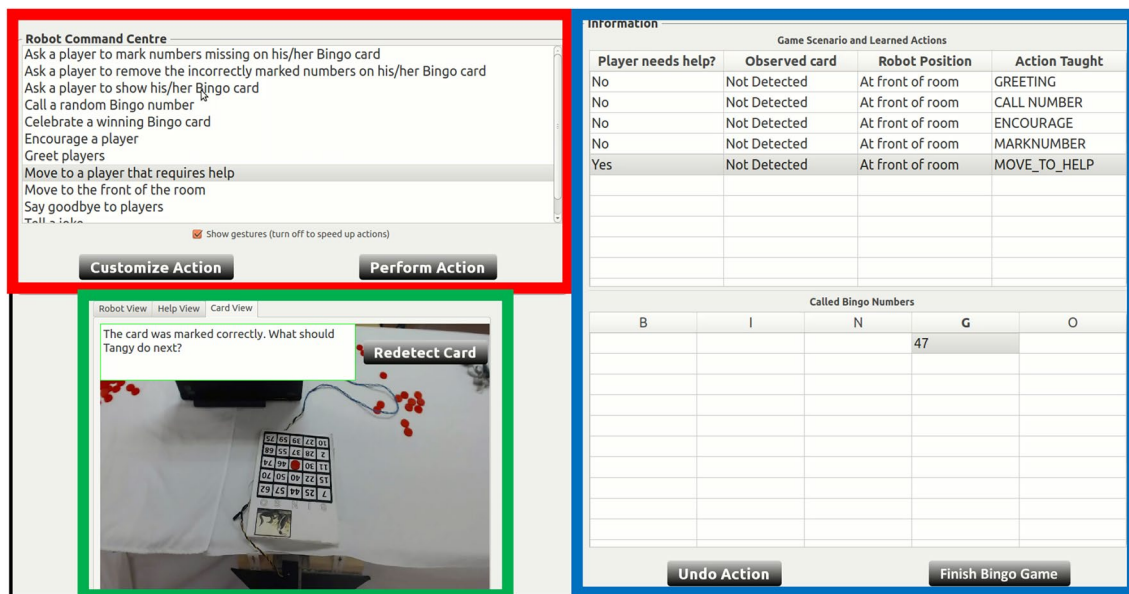


Fig. 3 Activity demonstration sub-interface consisting of: robot command centre (red), sensor views (green), and information section (blue). (Color figure online)

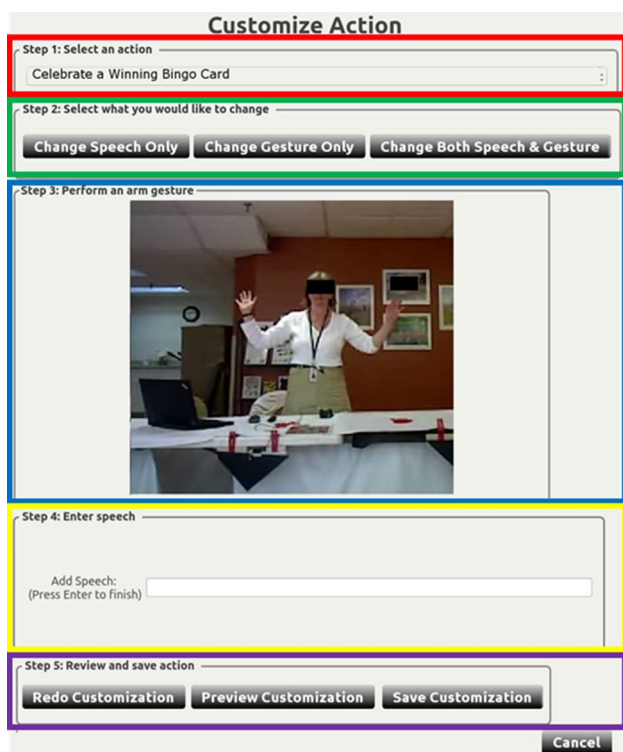


Fig. 4 Behavior demonstration sub-interface with the following steps: (1) selecting a behavior (red), (2) selecting what to change (green), (3) demonstrating an arm gesture (blue), (4) entering in speech (yellow), and (5) reviewing and completing a behavior demonstration (purple). (Color figure online)

during all steps of the learning process to improve the usability of the system:

Step 1 (red box) The teacher selects the behavior he/she would like to customize.

Step 2 (green box) The teacher is prompted to choose whether he/she would like to customize the speech and/or gestures. If a teacher has chosen to customize the gesture or both the speech and gesture, he/she proceeds to Step 3. Otherwise, if the teacher has chosen to only customize the speech for a behavior, he/she proceeds to step 4.

Step 3 (blue box) The teacher is prompted by the GUI to perform arm gestures for Tangy to display for the behavior.

Step 4 (yellow box) The teacher is prompted by the GUI to enter text he/she would like Tangy to say for the behavior.

Step 5 (purple box) The teacher can choose to preview, redo, or save a customization. When the teacher chooses to preview a behavior, Tangy performs the new speech and/or gesture that has been taught. If the teacher has chosen to redo a customization he/she returns to step 1. If the teacher has chosen to save the customization, then he/she returns to the activity demonstration sub-interface.

4.4 Identification of World State Parameters

In this work, we utilize the sub-modules we have developed in [12] to identify the world state parameters.

4.4.1 Person Identification

Player identities are determined using the OKAO™ Vision software library based on facial features from 2D images captured using the 2D Axis camera. The facial features are compared to a database of features of known players in order to determine the identity of a player.

4.4.2 Player Activity State

Player activity states are determined using images captured from the 2D Logitech camera. Namely, the unique symbol on a Bingo card is identified by computing the Speeded-Up Robust Features (SURF) on the images and matching these features to a database containing the SURF features for the unique symbol on each card. Once the symbol has been identified, the grid lines surrounding the numbers are recognized using a Hough transformation-based method. The location of each red circular marker is determined using a red blob filter. Identified red markers are then matched with their nearest neighbor grid squares. Squares that have been matched with a red blob are considered marked squares. The numbers in the marked squares are then compared to the set of numbers called by Tangy to determine the player activity state.

4.4.3 Help State

Player assistance requests are monitored using IR and 3D point cloud information from the Asus Xtion IR sensor behind Tangy and a Hough Transformation based methodology to identify IR reflective triangles that are revealed when a player presses the button on an assistance request device. The location of the player requesting for assistance is then determined by identifying the position of the IR triangle in the 3D point cloud of the environment.

4.4.4 Robot State

We use the ROS navigation software package for the robot to map, localize, and navigate in the environment using its wheel encoders and Hokuyo laser range finder. A map of the activity room is generated using the Gmapping Simultaneous Localization and Mapping technique. Robot localization is achieved using a probabilistic adaptive Monte Carlo localization approach which uses a particle filter to determine the robot's position. The A* pathfinding algorithm is utilized to plan a global path for the robot to follow and a

Dynamic Window based local path planner is responsible for implementing the trajectories.

5 Robot Teaching Study

We conducted a preliminary robot teaching study with caregivers at a residential care facility to investigate user experience and workload, and the usability of the robot *Lfd* system (Fig. 5).

5.1 Participants

Five caregivers ages 26–52 years old ($\mu=37$, $\sigma=9.80$) participated in a one-hour robot teaching session. They were either recreation programs staff or social services staff, and all were experienced with facilitating Bingo games with residents. Informed written consent was obtained prior to the study. Participants had no prior experience working with robots. Similar sample sizes ranging from 3 to 6 participants have been utilized in studies where the perceptions of domain specialists (e.g. elementary school teachers, informal older adult caregivers, therapists for autism) are investigated towards end-user customization, programming, or teaching interfaces [13, 25–27]. Smaller sample sizes are often utilized because there are challenges with access to specialist populations due to the small number of specialists in each facility and their limited availability [26]. Although these sample sizes are small it has been well established within the field of human factors that a sample of five participants is sufficient for identifying the usability problems of a system with diminishing returns with each additional participant [28, 29].




Fig. 5 Caregivers teaching the Bingo activity to Tangy by: **a** teaching the robot to request a player to bring the Bingo card forward when the robot cannot see the card, and **b** teaching the robot to wave for a Bingo behavior

5.2 Methods

Robot teaching sessions were conducted in an activity room in the residential care facility. Each robot teaching session began with Tangy at the front of the room and the teacher seated behind a table with a Bingo card, an assistance request device, and red circular markers. The teacher had access to the integrated demonstration GUI on a laptop.

A member of the research team conducted a 15-minute interactive tutorial with each participant on how to utilize the activity and behavior demonstration sub-interfaces to teach Tangy an activity by guiding them through the available features. Then, each participant was asked to use the system to teach Tangy a complete Bingo game so that all states and possible behaviors could be considered. Participants had access to all the default behaviors presented in Table 1 but did not have access to any prior demonstrations by other caregivers. A post-interaction questionnaire was administered to the participants after the teaching sessions were completed.

5.3 Measures

We administered a three-part post-interaction questionnaire to measure: (1) user experience of the *Lfd* system and GUI; (2) perceived workload during the teaching task; and (3) perceived usability of the *Lfd* system.

Part A User Experience Open-ended questions were administered to investigate user experience of the demonstration learning GUI and the teaching interaction with Tangy. The questions focused on: (1) the usefulness of the types of information presented on the demonstration interface, (2) the appropriateness of the modalities to teach Tangy new behaviors, (3) the level of enjoyment while teaching the robot, (4) the overall perceptions of the usefulness of the system, and (5) alternative activities that can be taught using the system.

Part B Perceived Workload Perceived workload was measured to determine the demand on a teacher while teaching the robot. This is important as our goal is to reduce the overall workload on caregivers in residential care facilities. We measure perceptions of workload utilizing the NASA-TLX task load index [30].

Part C Perceived Usability Perceived usability is measured in order to determine if caregivers can effectively and efficiently interact with the system to accomplish the robot teaching task. We utilize the System Usability Scale (SUS) [31] to measure perceptions of the usability of the *Lfd* system because it is a standard questionnaire which has often been utilized with end-users for end-user teaching or programming interfaces for robots [13, 32–34].

5.4 Results and Discussion

5.4.1 User Experience

We analyzed the responses to the open-ended questions by conducting a thematic analysis on the responses.

Interaction Modalities All five participants found it easy to demonstrate behaviors to Tangy using the modalities available to them. All participants agreed that demonstrating arm gestures was easy and commented on its usefulness for recreational activities. This aligns with our original design intent. Namely, Tangy learning behaviors from external observations allows a teacher to easily generate natural looking motion trajectories [19]. Participants 1 and 5 suggested they would also be comfortable teaching gestures by direct teleoperation of Tangy's arms. However, such a technique can be especially difficult if a user needs to generate gestures requiring both arms. This can lead to unnatural looking motions as teachers often produce sharper velocity changes when simultaneously manipulating two arms [19].

All participants also found typing the robot's speech through the GUI was easy and user friendly. Participants 1 and 3 further elaborated on this with such supporting statements as "inputting speech on a computer is good and doesn't need to be changed", and "typing is user friendly".

Useful GUI Features Participants found that the various features on the teaching GUI were appropriate for teaching the robot. Participants 1, 3, 4 and 5 found that it was very helpful to have the various views and perceived world states presented on the GUI as it provided them with awareness from the robot's perspective. This validates our aim to ensure that the teacher's interpretation of the world state is aligned with Tangy's. Most of the participants also found it useful that the GUI presented the scenarios that had already been taught during the session. They stated that they used this information to maintain flow of the teaching session by recalling what scenarios had been taught and what should still be taught.

Enjoyability The participants all enjoyed teaching the robot and made statements such as, "I found it interesting to teach the robot". Furthermore, they wanted to use the *LfD* system again in the future and made statements such as "I am excited to come back and use the system".

Ease of Use All the participants found the *LfD* system easy to use. Participants 1 and 2 further indicated that they would like more teaching trials to better familiarize themselves and become more comfortable with the system. Participant 1 elaborated by saying, "I am not very tech savvy, but am encouraged to come back and learn to use the system".

Teaching Speed It took on average 20 min ($\sigma=4.62$) for a caregiver to teach the robot. Based on the responses to the open-ended questions, the caregivers could envision "setting

aside time to teach a new activity" and stated ideally "the teaching would only need to take 15 min to teach". Participants 2 and 4 explicitly mentioned that they felt the teaching task could be sped up. Participant 4 elaborated by saying that "the robot should move faster during teaching. However, [the robot behaviors] do not need to be sped up for residents".

Currently, the angular joint movement speed for the robot's arms is limited to 15 deg/s and robot navigation speed is limited to 0.5 m/s for interaction with residents. In the future, we will investigate increasing these speeds for faster robot teaching, but at the same time ensuring that they are safe for the caregivers. We limited the robot's movements to be slower than typical human speeds because studies have shown that in general humans interacting with a mobile service robot prefer that it move slower than a human with speeds between 0.4 and 1 m/s being most acceptable [35]. This was confirmed by the caregiver's response to our open-ended questions specifying that the speed was appropriate for residents. Furthermore, residents did not have any negative opinions towards the speed of the robot during the activity facilitation HRI study detailed in Sect. 6.

It is also important to note that retraining would only be necessary in scenarios where caregivers wanted to: (1) create a new behavior and/or make customizations to the robot's existing behaviors, or (2) teach the robot new activity scenarios previously unobserved during the demonstration. Herein, it would not be necessary for the teachers to demonstrate the entire activity again. Instead, individual behaviors can be created or customized using the demonstration interface. Similarly, previously unobserved activity scenarios can be taught to the robot by having a teacher setup the desired scenario and selecting/creating an appropriate behavior for the robot to implement using the activity demonstration interface. These activity scenario demonstrations would then be added as additional world state-behavior pairs to the existing activity trajectory and used as input into the activity learner to learn an activity policy.

Importance and Usefulness of Behavior Customization All participants emphasized that the ability to customize robot actions was important and useful. Participant 2 further elaborated saying that "customizing speech is good for making the game exciting [for residents]". Participant 5 mentioned that "residents ask a lot of questions in the real game, so customizing the robot's speech to explain the game is important". Furthermore, Participant 1 mentioned that "[modifying speech] is very useful and a great option".

Additional Activities All the participants stated that they would want to develop new recreational activities for Tangy using the *LfD* system. These activities included trivia, charades, and picture therapy games. The *LfD* system can be extended to include the aforementioned activities. The current version of the proposed learning from demonstration system learns from non-experts the high-level structure of a

multi-stepped social activity and the individual robot behaviors necessary to facilitate the activity. However, for full autonomy of all modules within the architecture, the system would need to include the detection and classification of the activity states as well. Currently, the state identification methods for monitoring the specific world state for an activity (e.g. Bingo card state and assistance requests) using the robot’s sensors have been designed by the researchers. In the future, relevant state information for the facilitation of an activity could be directly taught by caregivers. This could be accomplished by extending our proposed system architecture to incorporate a similar approach to that presented in [36] which has been shown to enable human users to teach a robot to classify object and environmental states given the robot’s sensory information. Namely, the approach had users interactively train a robot by demonstrating the state of the environment to the robot and providing a label for the demonstrated state. The robot therefore learns to determine the current state given specific input sensory information.

5.4.2 Perceived Workload

The NASA-TLX questionnaire results for our system are presented in Table 2 and Fig. 6. The results showed moderate perceived mental ($\mu = 11, \sigma = 0.39$) and temporal ($\mu = 10.2, \sigma = 0.86$) demand during the sessions. The teaching task requires a teacher to design the sequence of scenarios and behaviors he/she would like for facilitating a complete game which is directly linked to mental demand. The moderate perceived temporal demand is also expected as caregivers had to facilitate a complete game. Some of the participants who scored higher for temporal demand were concerned with their general technology aptitude and mentioned that teaching may become faster as they become more familiar with the system. Perceived physical demand ($\mu = 4, \sigma = 0.95$) was low and perceived




Fig. 6 Participant NASA-TLX scores during the activity teaching task

effort ($\mu = 6.6, \sigma = 1.15$) and frustration ($\mu = 5.6, \sigma = 0.90$) were moderately low during the teaching task. The physical activity that occurred was only using the mouse and keyboard or demonstrating gestures with their arms. The participants’ scores for effort and frustration are supported by their feedback from the open-ended questions towards the ease of use of the interface. Participant 4 was moderately frustrated during the teaching task, which can be supported by this participant’s suggestion in the open-ended questions towards having the robot move faster when navigating and performing gestures. Overall, the participants felt they performed well during the teaching task and had moderately good perceived performance scores ($\mu = 5.8, \sigma = 0.66$). The NASA-TLX workload scores ranged from 34.67 to 51.33 ($\mu = 41.4, \sigma = 5.53$).

According to a meta-analysis of over 200 publications utilizing the NASA-TLX [37], daily activity tasks had a median of 18.30 for their reported mean global workload scores. Tasks such as operating a robot ($\bar{x} = 56$) had a higher median for their reported mean global workload scores. Our mean global workload score is within the

Table 2 NASA-TLX Questionnaire Scores

Workload factor	Participants										Mean Score	SD
	1		2		3		4		5			
	Weight	Score	Weight	Score	Weight	Score	Weight	Score	Weight	Score		
Mental demand	0.20	10	0.27	10	0.2	14	0.27	10	0.20	11	11.0	1.55
Physical demand	0	1	0.07	10	0.1	1	0	1	0	7	4.0	3.79
Temporal demand	0.20	11	0.27	12	0.20	15	0.27	8	0.07	5	10.2	3.43
Performance	0.20	3	0.20	10	0.27	3	0.20	6	0.33	7	5.8	2.64
Effort	0.33	10	0.13	10	0.13	1	0.13	1	0.20	11	6.6	4.59
Frustration	0.07	4	0.07	6	0.07	1	0.13	12	0.20	5	5.6	3.61
Overall workload	42.00		51.33		34.67		38.70		40.33		41.4	5.53
Objective workload: total no. of robot behaviors and world states created	86		97		87		77		66		82.6	11.67

lower quartile of mean global workload scores for the robot operation tasks reported in [37].

In Table 2 we also provide objective data on participants' workload during demonstrations. Their objective workload is defined by the total number of behaviors and world states they created. Participants on average created 82.6 ($\sigma = 11.67$) behaviors and world states. A Pearson's correlation coefficient was utilized to identify if there was a relationship between their subjective and objective workloads. We utilized an $\alpha = 0.05$ for our analyses. The results of the correlation analysis between subjective and objective workload showed a Pearson correlation of $r = 0.514$ with no significant relationship ($p = 0.376$). Hence, these results demonstrate that there was no statistically significant relationship between participants' subjective evaluation of workload and objective measures of their workload. Although we did not observe a statistically significant relationship, it is interesting to note that there were only small variations in the subjective workload scores and the objective workload values across all the participants. This suggests that the perceived and objective workload between participants was consistent when using the *LfD* system.

5.4.3 Perceived Usability

The SUS scores ranged from 42.5 to 75 ($\mu = 58$, $m = 60$, $\sigma = 11.11$). The mean SUS scores can be interpreted with an *OK* adjective rating, which suggests that there is some room

for improvement with respect to the usability of the system [38]. As suggested in the open-ended questions, improvements could include providing more teaching trials to allow caregivers to become more comfortable with the system and increasing the teaching speed.

6 HRI Study with Residents

An HRI study with residents was conducted to: (1) evaluate the ability of the system to learn and facilitate Bingo games from the teacher demonstrations and (2) investigate residents' experience and perceptions of Tangy facilitating the learned activity for them.

6.1 Methods

A total of 404 world state-behavior pairs were observed from the demonstrations provided by the five caregivers with an average of 81 behaviors per caregiver. Table 3 below provides the frequency each world state-behavior pair was observed during the demonstrations by the five caregivers and Fig. 7 illustrates the behaviors they utilized. As noted in Table 3 different behaviors were selected by the caregivers for the same world state. A video of the example robot behaviors during the Bingo game can also be viewed at <https://youtu.be/wmmsza9QVTg>. We verified that only a single demonstration would be necessary for the random forest

Table 3 Observed frequency of each world state during caregiver demonstrations

Current robot behavior	Robot location	Help state	User activity state	Demonstrator selected action	Total instances of world state
None	Front of room	ANR	Occluded	Greet players	5
Call out a Bingo number	Front of room	ANR	Occluded	Encourage a player	49
Navigate to the front of the room/ greet players/call out a Bingo number/ tell a joke/encourage a player	Front of room	ANR	Occluded	Call out a Bingo number	271
Call out a Bingo number	Front of room	ANR	Occluded	Tell a joke	10
Navigate to player	At user	AR	Bingo	Celebrate a winning Bingo card	5
Navigate to player	At user	AR	Incorrectly Marked	Ask a player to remove incorrectly marked numbers on Bingo card	3
Navigate to player	At user	AR	Missing Numbers	Ask a player to mark missing numbers on Bingo card	11
Navigate to player	At user	AR	Occluded	Ask a player to show Bingo card	3
Call out a Bingo number/tell a joke/ encourage a player	At front of room	AR	Occluded	Navigate to a player who requires help	22
Ask a player to mark missing numbers on card/ celebrate a winning Bingo card/ ask a player to remove incorrectly marked numbers on Bingo card	At user	ANR	Occluded	Navigate to the front of the room	16
Celebrate a winning Bingo card	At user	AR	Bingo	Say goodbye to players	5

ANR assistance not required, *AR* assistance required




Fig. 7 Robot behaviors during a facilitated Bingo game with users: **a** greeting; **b** call number; **c** joke; **d** navigate; **e** request to remove markers from numbers that have not been called; **f** request to move card

closer to robot; **g** request to mark numbers that have been called; **h** encourage user to keep up the good work; and **i** celebrate

to converge to an accurate policy for a Bingo game if all the potential scenarios during a Bingo game were demonstrated by a caregiver at least twice during a session. In our robot teaching study, each caregiver did not demonstrate all the possible scenarios in a Bingo game before the sessions ended and each only accounted for a portion of the potential scenarios even though they demonstrated a complete Bingo game (i.e. winning Bingo card condition). Hence, we combined the caregivers’ demonstrations into a single dataset and used it as input into the activity learning sub-module

to learn a Bingo activity policy that models the combined facilitation behaviors of all the caregivers. An example of a decision tree from the learned random forest is presented in Fig. 8. Since this was the caregivers first interaction with the *LfD* system they were likely more focused on learning the interface and completing a Bingo game then demonstrating all the potential activity scenarios. We hypothesize that given more time and familiarity with the *LfD* system the caregivers would provide a more complete demonstration of all the possible activity scenarios during a game.

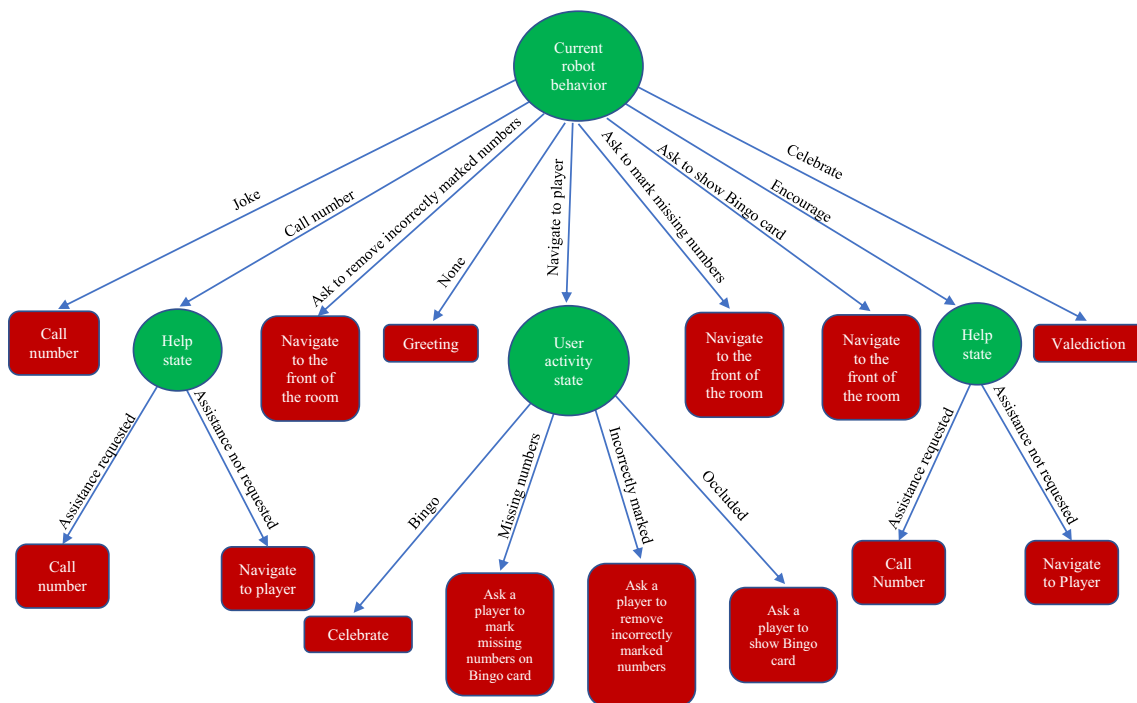


Fig. 8 Example decision tree from the random forest learned from the caregivers’ demonstrations. Decision nodes (green circles) represent the state being evaluated. Branches (blue edges) represent the pos-

sible values for a state. Leaf nodes (red rectangles) are the executed actions. (Color figure online)

With its new learned activity policy and behavior policies, Tangy facilitated Bingo games with groups of residents that were cognitively intact or had mild cognitive impairment according to the Cognitive Performance Scale, Fig. 9. A total of 18 participants (2 Males, 16 Females) ranging in age from 73 to 93 ($\mu = 83.06$, $\sigma = 6.27$) participated in one-hour Bingo sessions with the robot. Post-sessions, each participant completed a short 5-point Likert scale questionnaire (5- strongly agree, 1-strongly disagree) based on their experience with Tangy and the Bingo game, Table 4. Similar samples sizes ranging from 3 to 26 participants have been utilized in HRI studies with older adults [39].

6.2 Results and Discussions

Tangy was capable of facilitating the Bingo games with the residents with a 100% success rate using the learned activity policy. This demonstrates that non-expert caregivers can teach a high-level activity structure to a robot and the robot can learn an optimal policy to autonomously implement the activity.

Table 4 presents the descriptive statistics for the questionnaire results from this study and from our prior work in



Fig. 9 A Bingo game facilitated by Tangy with a group of residents

[40] where a traditional finite state machine developed by our research team was used to have Tangy autonomously facilitate Bingo games with older adults at a long-term care facility. In our prior work, the finite state machine was customized for the facility based on our observations of the structure of a Bingo game facilitated by caregivers at the facility and their individual activity facilitation behaviors. Overall the questionnaire results demonstrated that the residents had positive perceptions towards the robot facilitated Bingo games. Namely, the results suggest that, in general, the residents had positive and engaging experiences because they found the robot facilitated Bingo game interesting ($\bar{x} = 4$); enjoyed playing the Bingo game with Tangy ($\bar{x} = 5$); and found the robot was able to help them during the Bingo game ($\bar{x} = 4.5$). The participants also agreed that Tangy should host Bingo games again ($\bar{x} = 5$) and intended to participate in future robot facilitated games ($\bar{x} = 5$). These results demonstrate that the proposed system architecture has the potential for enabling caregivers to effectively teach a robot to facilitate group recreational activities that are engaging to residents and would lead to future interactions.

The questionnaire results from this study are similar to the questionnaire results with the Bingo games facilitated by the designed finite state machine. This demonstrates that our system would enable non-expert caregivers to independently teach a robot a customized activity with customized behaviors for their specific facility and the learned policy would perform just as well as one customized by an expert roboticist based on direct observations of a caregiver. However, pre-programming a robot to autonomously facilitate customized activities requires a considerable amount of expertise and is not scalable as it would require robotics professionals to observe caregivers at each individual facility and customize activities according to these observations. Hence, the

Table 4 Comparison of questionnaire results with [40]

Statement	Median	Frequency					Median from [40]	Frequency from [40]				
		1	2	3	4	5		1	2	3	4	5
I think Tangy makes the Bingo game interesting	4	2	2	1	5	8	5	0	0	0	0	7
I enjoyed playing Bingo with Tangy	5	0	1	1	5	11	5	0	0	0	0	7
Tangy was able to call out the Bingo numbers clearly.	5	0	0	0	0	18	N/A	-	-	-	-	-
Tangy displaying numbers on his screen was helpful	5	1	0	1	0	16	N/A	-	-	-	-	-
I enjoyed when the robot waved to me	5	0	0	5	3	10	N/A	-	-	-	-	-
I thought it was helpful when the robot pointed to its screen	3.5	1	0	8	2	7	N/A	-	-	-	-	-
I enjoyed when robot did the celebration gesture	5	0	0	1	6	11	N/A	-	-	-	-	-
Tangy was able to help me	4.5	1	1	3	4	9	5	0	0	1	1	5
I would ask Tangy for help again	5	1	1	0	3	13	5	0	0	1	0	6
I would follow what Tangy asks me to do in the Bingo game	5	0	0	0	3	15	5	0	0	0	0	7
I would play Bingo with Tangy again.	5	1	0	0	5	12	5	0	0	0	0	7
Tangy should host Bingo games again	5	0	0	2	5	11	5	0	0	0	1	6

main advantage of the proposed approach is that it would not limit a socially assistive robot to only include activities pre-programmed on the robot by roboticists but would enable caregivers to create and customize their own activities.

It is interesting to note that the caregivers taught a Bingo activity policy that requires Tangy to frequently ask for player feedback (e.g. asking how everyone is doing during the game) and provide instructions (e.g. letting players know the number pattern they need to mark on their cards to win) to all the players participating in the activity. This was achieved by creating new speech for the robot's encourage behavior and implementing this behavior when the robot was facilitating the activity at the front of the room. This is different from our original intentions of encouraging a player when the robot identifies a correctly marked card [12]. This is important to note, as this behavior is often utilized by caregivers to actively encourage all the older adults participating in an activity regardless of their activity performance.

Our expectation that caregivers would execute multiple different behaviors in the same world state during the facilitation of a Bingo activity was also confirmed by the learned activity policy. Namely, multiple behaviors were demonstrated for the same world state of the robot at the front of the room and when players have not requested for assistance. In particular, Tangy learned a policy where in this world state there was a probability of 82% it would call Bingo numbers, a probability of 15% it would encourage players to check their cards, and a 3% probability it would tell jokes.

Both the actual use of the customization capabilities of the *LfD* system by caregivers as well as their responses towards the usefulness and importance of behavior customization supports our motivation of adapting robot facilitated activities based on the needs of the users within the residential care facilities. The positive feedback from the older residents provides further insight into the benefits of robot activity personalization by non-expert teachers.

7 Conclusions

We have presented a *LfD* system architecture for SARs to address the unique challenges of learning from non-experts group social activities. Namely, we proposed a random forest classifier which learns a probabilistic model of a human teacher's non-deterministic activity facilitation behavior from his/her demonstrations and implements the learned activity where the world state is dynamic (i.e. the behaviors of the robot cannot be sequentially planned). We evaluated the user experience, perceived workload, and perceived usability of the *LfD* system by having caregivers from a residential care facility teach Tangy a Bingo activity. Results showed that, overall, the caregivers had a positive experience using the system and would want to teach the robot

to autonomously facilitate various activities with residents. They had moderately low perceived workloads while utilizing the system and found it easy to use as well as useful. Furthermore, the caregivers taught the robot to execute multiple different behaviors in the same world state. This validates our expectation that caregivers want to demonstrate non-deterministic behaviors to accommodate the residents' needs and our proposed *LfD* system could learn a probabilistic model of the caregivers' non-deterministic activity facilitation behavior from their demonstrations. The importance of caregivers being able to personalize the activity was also emphasized by the older residents as they found the personalized behaviors helpful as well as entertaining and enjoyed playing Bingo with Tangy.

The presented proof-of-concept study consisted of short-term interactions between the residents and the robot. Hence, we will conduct more long-term studies to investigate if there is a novelty effect with the robot. Furthermore, participant perceptions of the technology were obtained via self-reported surveys so there could be potential for response bias. Based on the results we have obtained here, in the future, to address these limitations we will run long-term studies with both caregivers and residents at residential care facilities to observe their continued use of the robot and measure participants' pre- and post-study perceptions towards the technology to investigate any changes in perceptions over time.

A limitation of the implemented *LfD* system is that it currently assumes players will be engaged in the group recreational activity facilitated by the robot (e.g. playing the game) using the activity policy learned from the caregivers. Although all the residents during our HRI study were engaged during our Bingo games, the activity policy taught by the caregivers may not be the optimal policy for keeping residents engaged. The proposed framework could be modified in the future to learn an optimal policy on-line through a two-step process. First, the system could be developed to monitor residents' engagement in an activity from their facial expressions, body language, eye contact, and verbal behaviors. Resident engagement can then be utilized by the socially assistive robot as feedback to identify robot behaviors which maintain and/or improve resident engagement during an activity. Namely, the socially assistive robot can utilize resident engagement as the reward in on-line reinforcement learning based techniques to learn the optimal policy for maintaining engagement during an activity.

Although both caregivers and residents from the residential care facility had positive perceptions towards the robotic technology, in the future we intend to conduct a comparative analysis between traditional human facilitated activities and the robot facilitated activities. Such a study will provide insights on how the proposed *LfD* system and the robot facilitated activities could be utilized to support residential care facilities in providing meaningful activities to residents.

Acknowledgements The authors would like to thank the staff and residents of our partner care facility for their support, assistance, and participation in these user studies. The authors also thank Sharaf Mohamed, Francis Despond, Christina Moro, and Tan Zhang for their assistance with the robot and teaching study.

Funding This study was funded by the Natural Sciences and Engineering Council of Canada (NSERC), Dr. Robot Inc. and the Canada Research Chairs (CRC) program.

Compliance with Ethical Standards

Conflict of interest The authors declare that they have no conflict of interest.

Informed Consent Written consent was obtained from all participants prior to their participation.

Appendix


Post-Interaction Questionnaire

Part A: Open-Ended Questions

1. Did you find it easy to customize a new gesture for the robot?
2. Would you prefer another method for teaching the robot a new gesture (e.g. keyboard and mouse, physically moving the robot's arms)?
3. Did you find it easy to customize new speech for an action?
4. Would you prefer another method for inputting new speech (e.g. such as speaking into a microphone)?
5. Was it useful to see what the robot saw in the Graphical User Interface?
6. Was it useful to present information about the current activity that the robot is observing (e.g. if someone has requested for assistance, how a Bingo card has been marked by a player, the called out Bingo numbers).
7. Was it useful to know what the robot has learned during the robot teaching session in the Graphical User Interface?
8. Do you think you could develop new recreational activities for the robot to do with the residents using this system?
9. Did you enjoy teaching the robot a new activity?
10. Do you have any additional comments or suggestions?

Part B: NASA-TLX Task Load Index


Please place an "X" along each scale at the point that best indicates your experience during the robot teaching session, ranging from low to high for statements 1-5 and good to bad for statement 6.



1. *Mental Demand*: How much mental and perceptual activity was required during the Bingo teaching task (e.g. moving arms, clicking the mouse, pressing buttons)? For example, was the Bingo teaching task easy or demanding, simple or complex, exacting or forgiving.
2. *Physical Demand*: How much physical activity was required during the Bingo teaching task (e.g. pushing, pulling, turning, controlling, activating, etc.)? Was the Bingo teaching task easy or demanding, slow or brisk, slack or strenuous, restful or laborious?
3. *Temporal Demand*: How much time pressure did you feel due to the rate or pace at which the Bingo teaching task occurred? Was the pace slow and leisurely or rapid and frantic?
4. *Effort*: How hard mentally or physically did you have to work to teach the robot Bingo?
5. *Frustration*: How discouraged, stressed, irritated, and annoyed versus gratified, relaxed, content, and complacent did you feel while teaching the robot.
6. *Performance*: How successful do you think you were in accomplishing the goals of the Bingo teaching task? How satisfied were you with your performance in accomplishing these goals?

Part C: System Usability Scale

Please place an "X" to indicate your level of agreement to the following statements.



1. I think that I would like to use the robot teaching system frequently to teach the robot to facilitate new leisure activities with residents.
2. I found using the robot teaching system too complex.
3. I thought the robot teaching system was easy to use.
4. I think that I would need the support of a technical person who is always nearby to be able to use this robot teaching system.
5. I found the various functions of the robot teaching system were well integrated.

6. I thought there was too much inconsistency in the robot teaching system.
7. I would imagine that most recreational activity program staff would very quickly learn to use the robot teaching system.
8. I found the robot teaching system very cumbersome to use.
9. I felt very confident using the robot teaching system.
10. I needed to learn a lot of things before I could get going with the robot teaching system.

References

1. Thompson C et al (2017) The robot tangy facilitating trivia games: a teambased user-study with long-term care residents. In: IEEE international symposium on robotics and intelligent sensors, pp 173–178
2. Moyle W et al (2013) Exploring the effect of companion robots on emotional expression in older adults with dementia: a pilot randomized controlled trial. *J Gerontol Nurs* 39:46–53
3. Janke MC, Payne LL, Van Puymbroeck M (2008) The role of informal and formal leisure activities in the disablement process. *Int J Aging Hum Dev* 67:231–257
4. Fernández-Mayoralas G et al (2015) Active ageing and quality of life: factors associated with participation in leisure activities among institutionalized older adults, with and without dementia. *Aging Mental Health* 19:1031–1104
5. Kolanowski AM et al (2001) Capturing interests: therapeutic recreation activities for persons with dementia. *Therap Recreat J* 35:220–235
6. Chernova S, Thomaz AL (2014) Robot learning from human teachers. *Synth Lect Artif Intell Mach Learn* 8:1–121
7. Huang J, Lau T, Cakmak M (2016) Design and evaluation of a rapid programming system for service robots. In: ACM/IEEE international conference on human–robot interaction, pp 295–302
8. Niekum S et al (2015) Learning grounded finite-state representations from unstructured demonstrations. *Int J Robot Res* 34:131–157
9. Mohseni-Kabir A et al (2015) Interactive hierarchical task learning from a single demonstration. In: IEEE international conference on human–robot interaction, pp 205–212
10. Zhu Z, Hu H (2018) Robot learning from demonstration in robotic assembly: a survey. *Robotics* 7:1–25
11. Ryvicker M (2011) Staff-resident interaction in the nursing home: an ethnographic study of socio-economic disparities and community contexts. *J Aging Stud* 25:295–304
12. Louie W-YG, Nejat G (2016) A learning from demonstration system architecture for robots learning social group recreational activities. In: IEEE international conference on intelligent robots and systems, pp 808–814
13. Saunders J et al (2016) Teach me—show me?—End-user personalization of a smart home and companion robot. *IEEE Trans Hum Mach Syst* 46:27–40
14. Uyanik KF et al (2013) Learning social affordances and using them for planning. *Proc Ann Meet Cognit Sci Soc* 35:3604–3609
15. Louie WG et al (2014) A focus group study on the design considerations and impressions of a socially assistive robot for long-term care. In: IEEE international symposium on robot and human interactive communication, pp 237–242
16. Breiman L (2001) Random forests. *Mach Learn* 45:5–32
17. Fernández-Delgado M, Cernadas E, Barro S, Amorim D (2014) Do we need hundreds of classifiers to solve real world classification problems? *J Mach Learn Res* 15:3133–3181
18. Biau G, Scornet E (2016) A random forest guided tour. *Test* 25:197–227
19. Calinon S, Billard A (2007) Incremental learning of gestures by imitation in a humanoid robot. In: ACM/IEEE international conference on human–robot interaction, pp 255–262
20. Zhang T et al (2018) Robot imitation learning of social gestures with self-collision avoidance using a 3D sensor. *Sensors* 18:1–17
21. Kunstler R, Daly FS (2010) Therapeutic recreation leadership and programming. Human Kinetics Publishers, Champaign
22. Smith MC et al (1994) Age and skill differences in adaptive competence. *Int J Aging Hum Dev* 39:121–136
23. Sobel BP (2001) Bingo vs. physical intervention in stimulating short-term cognition in Alzheimer’s disease patients. *Am J Alzheimers Dis Other Dement* 16:115–120
24. Tak SH et al (2014) Activity engagement: perspectives from nursing home residents with Dementia. *Educ Gerontol* 41:182–192
25. Griol D, Callejas Z (2013) An architecture to develop multimodal educative applications with Chatbots. *Int J Adv Rob Syst* 10:1–15
26. Buchina N, Kamel S, Barakova E (2016) Design and evaluation of an end-user friendly tool for robot programming. In: IEEE international symposium on robot and human interactive communication, pp 185–191
27. Barakova EI, Gillesen JC, Huskens BE, Lourens T (2013) End-user programming architecture facilitates the uptake of robots in social therapies. *Robot Auton Syst* 61:704–713
28. Virzi RA (1992) Refining the test phase of usability evaluation: how many subjects is enough? *Hum Factors* 34:457–468
29. Turner CW, Lewis JR, Nielsen J (2006) Determining usability test sample size. *Int Encycl Ergon Hum Fact* 3:3084–3088
30. Hart G, Staveland LE (1988) Development of NASA-TLX (Task Load Index): results of empirical and theoretical research. *Adv Psychol* 52:139–183
31. Brooke J et al (1996) SUS—a quick and dirty usability scale. *Usability Eval Ind* 189:4–7
32. Kronander K, Billard A (2014) Learning compliant manipulation through kinesthetic and tactile human–robot interaction. *IEEE Trans Haptics* 7:367–380
33. Weiss A et al (2009) Teaching a humanoid: a user study on learning by demonstration with HOAP-3. In: IEEE international symposium on robot and human interactive communication, pp 147–152
34. Blume J et al (2013) Programming concept for an industrial HRI packaging cell. In: 2013 IEEE international symposium on robot and human interactive communication, pp 93–98
35. Butler JT, Agah A (2001) Psychological effects of behavior patterns of a mobile personal robot. *Auton Robots* 10:185–202
36. Forbes M et al (2014) Grounding antonym adjective pairs through interaction. In: ACM/IEEE international conference on human-robot interaction—workshop on humans and robots in asymmetric interactions, pp 1–4
37. Grier R (2015) How high is high? A meta-analysis of NASA-TLX global workload scores. In: Proceedings of the human factors and ergonomics society, pp 1727–1731
38. Bangor A et al (2009) Determining what individual SUS scores mean: adding an adjective rating scale. *J Usability Stud* 4:114–123
39. Bemelmans R, Gelderblom GJ, Jonker P, De Witte L (2012) Socially assistive robots in elderly care: a systematic review into effects and effectiveness. *J Am Med Dir Assoc* 13:114–120
40. Li J, Louie W-YG, Mohamed S, Despond F, Nejat G (2016) A user-study with Tangy the Bingo facilitating robot and long-term care residents. In: IEEE international symposium on robotics and intelligent sensors, pp 109–115

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Wing-Yue Geoffrey Louie received his B.A.Sc. and Ph.D. degrees in mechanical engineering from the University of Toronto, Toronto, ON, Canada, in 2011 and 2017, respectively. In 2017, he joined the Department of Electrical and Computer Engineering as an Assistant Professor at Oakland University, Rochester, MI. He is presently the Director of the Intelligent Robotics Laboratory. His research focuses on integration of multi-sensor systems on robotic platforms, development of perception algorithms, human-robot/human-machine interfaces, high & low-level robot controller design, general robot learning, artificial intelligence, and the study of human-robot interactions. He applies his research to real-world applications scenarios including manufacturing, rehabilitation, service, and personal robotics as well as autonomous driving. One of his main research goals is the development of programmable social robots for everyone.

Goldie Nejat received her B.A.Sc. and Ph.D. degrees in mechanical engineering from the University of Toronto, Toronto, Canada, in 2001 and 2005, respectively. She is a Professor and the Director of the Autonomous Systems and Biomechanics Laboratory (ASBLab) in the Department of Mechanical and Industrial Engineering at the University of Toronto. She is also the Canada Research Chair in Robots for Society and an Adjunct Scientist at the Toronto Rehabilitation Institute. Her research interests include the development of intelligent assistive, social and service robots for a number of applications including healthcare, search and rescue, manufacturing and exploration/surveillance fields. One main aspect of her research is to develop and integrate intelligent socially assistive robots for assistive human–robot interactions with older adults, including those living with dementia, in healthcare facilities and private homes.