



Design and Validation of MOMDP Models for Child–Robot Interaction Within Tasks of Robot-Assisted ASD Diagnostic Protocol

Frano Petric¹ · Zdenko Kovacic¹

Accepted: 3 July 2019 / Published online: 23 July 2019
© Springer Nature B.V. 2019

Abstract

The existing procedures for autism spectrum disorder diagnosis are time-consuming and challenging both for evaluators and children being evaluated. Occurrence of low agreement rates between different clinicians when evaluating a child suggests that there exists a need for a more objective approach to diagnostics. To that end, we developed a robot-assisted ASD diagnostic protocol. In this work the focus is on robot reasoning for tasks of the protocol. We propose the mixed observability Markov decision process models for tasks which infer information about the state of a child based on observations of child's behavior. In order to formulate observation probabilities of task models, ASD experts are surveyed and their knowledge is encoded in the observation probabilities of task models. Expert knowledge also allowed for implementation of child behavioral models which are used to validate and tune developed models. Following the successful validation through simulations of child–robot interaction using child behavioral models, task models are validated through experimental sessions with six typically developing children and eight children with ASD. Results obtained through experiments show that the robot is capable of correctly identifying the behavior of the child within the diagnostic tasks.

Keywords Robotics · Autism spectrum disorder · Diagnostics · Mixed observability Markov decision processes

1 Introduction

The focus of this paper is the robot reasoning for robot-assisted autism spectrum disorder (ASD) diagnostics. Although ASD, a neurodevelopmental disorder with no medical markers, is becoming more commonly diagnosed [1–3], the diagnostic procedure for ASD has been identified as problematic, both in terms of inter-rater reliability [4] but also in terms of time needed to reach a diagnosis [5]. Direct and indirect costs linked to delay in the diagnosis [6,7] only

emphasize the already well-know importance of early diagnosis which facilitates early intervention. The problems with ASD diagnostics stem from the fact that the procedure is highly complex, consisting of simultaneous administration of various tasks and observation and coding of the behavior and requires substantial amount of training for a human examiner. Although the need for a more objective approach is recognized [8], according to recent surveys [9,10] which investigated the use of robots and technologies for people with ASD, there are no robots used in the diagnostic process. The most likely reason for such scarcity is the requirement for efficient processing and reasoning algorithms that are to be implemented on the robot to enable the robot to assist in the procedure without increasing the number of people needed. At the current stage of our work, we bypass those requirements by using a semi-autonomous robot in the robot-assisted ASD diagnostic protocol [11], which stands out as one of the first applications of humanoid robots in the diagnostics procedure but also adheres to the trends in the field of using robots for research related to ASD where most of the researchers use Wizard-of-Oz techniques [12–16]. The difference in our approach is in the fact that the robot is aided by providing correct observations of social cues in cases where

This work has been fully supported by the Croatian Science Foundation through the ADORE Project (HRZZ-93743-2014).

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s12369-019-00577-0>) contains supplementary material, which is available to authorized users.

✉ Frano Petric
frano.petric@fer.hr
Zdenko Kovacic
zdenko.kovacic@fer.hr

¹ Laboratory for Robotics and Intelligent Control Systems, University of Zagreb, Faculty of Electrical Engineering and Computing, Unska 3, Zagreb, Croatia

the robot fails to detect them automatically, and not remotely controlled. Operating under an assumption of near-perfect social cue detection, the goal of this work is to formulate a framework that will enable the robot to choose actions autonomously and more importantly to process the observations of child's behavior and infer information about the unobservable state of the child.

To that end, we are proposing a Mixed Observability Markov Decision Process (MOMDP) models for the tasks of the protocol. Our approach is similar to [17], where authors model the robot and mission states as fully observable, while the operator's cognitive ability is considered to be partially observable. Similar model is presented in [18], where authors model current robot configuration and history of the interaction as observable states, while the set of partially observable variables consists of human mode and human adaptability. A significant portion of research on using Partially Observable Markov Decision Processes (POMDP) and MOMDPs in robotic applications investigates methods to generate the knowledge representation by using various methods of machine learning on either already available data or data collected through experiments [19–21]. For a POMDP, this knowledge is for the most part encoded in form of state transition and observation probabilities and, to a lesser degree, reward functions. The design factors that we focus on in this work are transition and observation probabilities, while rewards are used to ensure the task structure is in accordance to the diagnostic protocol. To obtain aforementioned probabilities one could analyze the statistical data on how children react to different prompts, but to the best of our knowledge, such a dataset either does not exist or is not publicly available, especially for our target demographic of children aged 2–6. Data acquisition is restricted by the specific circumstances of reduced availability of children with autism. Therefore, to obtain the data to encode in our MOMDP models, we take the approach used to develop early expert systems and survey experienced ASD clinicians to provide their estimates of the rates in which some events in the diagnosis procedure occur, which we then encode into probabilities in the task models. Then we use the belief state of the model as an automatic evaluator of the child's behavior. We do not evaluate the MOMDP decision-making feature, as the structure of a task is fixed.

The contributions of this paper are twofold: (1) the tasks of the robot-assisted ASD diagnostic protocol are modeled as MOMDPs by encoding expert knowledge about the expected behavior of children; and (2) the diagnostic validity of the task models is confirmed through experimental sessions with fourteen children by comparing the robot's belief at the end of a task to assessment of the interaction provided by ASD experts.

The paper is organized as follows. Following the introductory section, Sect. 2 presents our prior work and other

information upon which our work is developed. In Sect. 3 we describe the design team, experimental setup and the methodology used to evaluate whether the robot can identify the behavior of a child. Section 4 brings forth the structure of task models, which is complemented by Sect. 5 which describes how the expert data is encoded in the aforementioned models. Section 6 shows how simulation was used to perform initial assessment of task models and tune some parameters of the models. Finally, the results of the experimental sessions are presented and discussed in Sect. 7 after which we conclude the paper by outlining guidelines for future work in Sect. 8.

2 Preliminaries

In the envisioned robot-assisted diagnostic protocol, the robot is actively participating in the diagnostic process, both through eliciting the interaction and observing the child's behavior.

2.1 Robot-Assisted ASD Diagnostic Protocol

The protocol, which went through several iterations [11, 22, 23], currently consists of four tasks developed upon the ADOS [24] protocol:

- The *response to a name call* (RNC) task of the protocol focuses on a child's ability to respond after being called by name, with the response classified as positive if eye contact is detected.
- The *joint attention* (JA) task of the protocol focuses on a child's ability to transfer attention from one robot to other robot, with the positive response being verified via eye contact with the other robot.
- The *play request: simultaneous multi-channel communication assessment* (PR) task is used to instigate the child's vocalizations and eye-contact in coordination with hand and body gestures, in order to assess the child's ability to communicate on multiple channels simultaneously.
- The goal of the *functional imitation* (FI) task is to evaluate the child's ability to imitate simple actions and consists of 2 imitation sub-tasks with two objects: a toy frog and a cup.

In all tasks, the robot first performs an action intended to provoke the response of the child. List of all available actions is presented in Table 1.

Table 2 summarizes which social cue is evaluated in which task of the robot-assisted ASD diagnostic protocol. Our focus and prior work on autonomously detecting eye contact, imitation, joint attention, gestures and verbal abilities of a child [25–27] is grounded in established research on relevance of social cues with respect to ASD: eye contact as a mea-

Table 1 Actions performed by the robot in the diagnostic protocol

Action	Channels	Description
Call child by name	Audio	Speech
Call child using a special reference	Audio	Speech
Transfer attention by turning head	Audio, motion	Speech and movement of the robots head
Transfer attention by pointing with hand	Audio, motion	Speech and movement of the robots head and arm
Attract attention of the child	Audio, visual, motion	Playing music, flashing LEDs and movement of the robots head
Play	Audio, visual, motion	Playing music, flashing LEDs and movement of the robots head, arms and legs
Demonstrate frog jumping gesture	Audio, motion	Speech instructions and movement of the robots arms
Demonstrate drinking from a cup gesture	Audio, motion	Speech instructions and movement of the robots arms

The communication channels for each action are presented

Table 2 Social cues tracked by the robot in tasks of the robot-assisted ASD diagnostic protocol

Task	Eye contact	Utterances	Imitation	Gestures
RNC	Yes	Yes	No	No
JA	Yes	Yes	No	No
PR	Yes	Yes	No	Yes
FI	Yes	Yes	Yes	No

sure of attention [28,29], verbal abilities of the child [30,31], imitation capabilities of the child and ability to use gestures [32,33]. Since the next action the robot should perform depends on the response of the child and on its unobservable inner state, the tasks of the protocol are suitable candidates to be modeled via Markov decision processes with unobservable states.

2.2 POMDPs and MOMDPs

A POMDP models an agent decision process in which the system is discrete and dynamics of the system are Markovian but the agent cannot directly observe the system state (i.e. the condition of a child). A POMDP is defined as a tuple (S, A, O, T, Ω, R) , where S is a set of hidden states of the system, A is a set of actions that can be performed, O is a set of observations generated by the system, T denotes the conditional transition probability between states depending on action $p(s'|s, a)$, Ω defines the conditional observation probability $p(o'|s', a)$, while R is a reward function. Since the state of the system is hidden to the decision maker, the decision maker maintains the belief state b as a probability distribution over S . Thus, $b(s)$ is the probability of a system being in the state $s \in S$. Based on the current value of belief b , the decision maker chooses the *best* action with respect to the expected reward.

A mixed observability Markov decision process [34] is a special case of POMDP, specifically, a factored POMDP model with mixed state variables. Fully observable state com-

ponents are represented as a single state variable X , while the partially observable components of system state are represented as a different state variable Y , factoring the state space of the model, $S = X \times Y$. Since each of the system state components has a corresponding state transition function, namely T_x and T_y , the complete MOMDP model is formally specified as a tuple $(X, Y, A, O, T_x, T_y, \Omega, R)$. Full observability of some state components can then be exploited for more efficient solving of the POMDP, as implemented in the SARSOP [35] solver, which is used in this work.

3 Methodology

In designing the robot-assisted protocol and interaction scenarios for the tasks of the protocol, we employ similar strategy to that reported in [36], connecting robotics engineers with experts in ASD diagnostics and rehabilitation. We also employ similar naming scheme for team members. Interaction scenarios are constantly tuned by the means of iterative development, which consists of rapid prototyping of basic behaviors, deploying the robot in the session with children and then updating the requirements for the robot performance. This approach is similar to that reported in [37]. The exploratory session presented at the end of this paper is the latest iteration of this development and provides guidelines for protocol POMDP reward design. The robots used in the robot-assisted ASD diagnostic protocol are NAO robots by Softbank Robotics, hardware version V4 running software version 2.1 [38].

3.1 Design Team

Our design team consists of three robotics engineers (REs) and three ASD experts (AEs). Backgrounds of team members are presented in Table 3. Role of the REs in the iterative development of the robot-assisted protocol is implementation of robot behaviors and cognitive abilities and control of the robot during sessions with children, while AEs provide

Table 3 Members of the design team

Member	Background	Education level	Expertise
RE1	Electrical engineering	PhD	Robotics, control
RE2	Electrical engineering	PhD	Robotics, control
RE3	Electrical engineering	PhD student	Robotics, human–robot interaction
AE1	Education and rehabilitation sciences	PhD	ASD diagnostics
AE2	Education and rehabilitation sciences	PhD	ASD diagnostics
AE3	Education and rehabilitation sciences	PhD	ASD intervention

relevant information in form of expected child behavior. AEs also perform initial assessment of the results of the sessions with children to provide guidelines for future development.

For the work presented in this paper, the role of REs is to formulate the POMDP/MOMDP models and process the data obtained from external experts and finally encode the obtained information in the probabilities of observations in developed models. On the other hand, AEs perform the selection of external experts and resolve ambiguous situations in the obtained dataset.

3.2 Survey Participant Selection

AEs sought experienced clinicians with ASD diagnosis and intervention background to complete the survey. Due to highly specific information that we needed to obtain and the requirement to have multiple years of clinical experience, only 8 participants performed the survey. Average stated clinical experience of those 8 participants is 8.125 years, min 3 years, max 14 years.

3.3 Data Collection and Analysis

The survey was administered through a simple Google form. Bulk of the questions were of type: *Given 10 children of type X, how many of them would react with Y to Your prompt Z*, where X denotes either typically developing child or a child with ASD, Y denotes observations in the tasks of the diagnostic protocol and Z denotes actions in the task of the protocol. Other questions pertained to duration and number of occurrences of eye contact, joint attention and expected number of actions to obtain a valid response (eye contact or correct imitation) for a given group of children which are aimed more towards our implementation of social cue detecting algorithm. Answers were exported from Google spreadsheet and parsed in Python, which was also used by REs to extract the observation probabilities. Data from sessions with children was collected through camera and microphones both on-board the robot and in the experimental room.

3.4 Experimental Setup and Participants

The sessions with children are performed in the Child Communication Research Laboratory at the Croatian Institute for Brain Research. The laboratory has a fully equipped examination and observation room, and is equipped with a two-way mirror to enable the observation and recording of the experiments. The recording system allows for video recording using a high-definition camera placed behind the two-way mirror and audio recording using an overhead microphone in the examination room, as shown in Fig. 1.

When executing diagnostic tasks, the first NAO robot is placed with its back to the mirror, while the other one is sitting on the floor to its left with all lights turned off except during the joint attention task when it activates to draw the attention of the child. The layout is set up in such a way that the child, if cooperating, should be facing the main robot and the camera for the majority of session, so that its reactions and facial expressions can be observed and recorded. The clinician is standing next to the main robot (on the right side of the table), observing the interaction between the robot and the child. The task of the clinician consists of providing comfort and security to both the child and the parent, but also to assist the robot during the imitation task by placing the object into the robot's hand and using tactile sensors on the head to signal that the robot should start the demonstration. This procedure is used to eliminate the possibility of failed grasps of the object and ensures that the imitation is always performed. A parent or a caregiver is also present in the room, usually sitting on the couch behind the child, to provide additional comfort and security for the child.

Behind the two-way mirror is a control and observation room where another clinician and a robot operator are observing the experiment. All experiments are recorded both by the robot's on-board camera and by an external camera placed behind the two-way mirror. Audio from the session is recorded by an overhead microphone and the microphones of the robot. The main robot is operating in a semi-autonomous mode, since instead of directly controlling the robot actions, the operator is aiding the robot by providing the correct observations which are critical for task and protocol execution,



Fig. 1 On the left, the layout of the room where sessions with children were taking place. In the middle, a snapshot of an experimental session during the functional imitation task. On the right, a snapshot from the joint attention task. (Color figure online)

as the algorithms for social cue detection are not accurate enough for fully autonomous execution. The critical observations are those that directly influence the execution of a task, i.e. if the child correctly imitates the gesture, the robot should not repeat it so it is critical to detect that the gesture has been imitated, but it is not critical to detect whether the child was speaking at the same time. It is important to note that for the evaluation of robot decision-making, all observations will be critical, but we are not at that point yet in our research. Other correct observations are obtained offline through video analysis after the session. For the joint attention task, the main robot sends commands to the second robot to activate when needed, while the second robot signals to the main if the child responded by detecting eye contact.

The operator also has the ability to pause execution at any point if the robot system fails (i.e. the robot falls down). In such cases, the clinician in the room is tasked with setting the robot into a safe starting position to enable soft restart of the session. However, in cases where such failure occurred, the children were reluctant to continue so the session was stopped. It is important to reiterate that the operator cannot influence the decision-making of the robot but only provides decision support by making sure that the observation the controller is receiving is indeed the correct one. Such approach is necessary in order not to waste a session due to incorrect critical observation detections. Within the task the robot chooses actions according to solutions of each task model. Task models are running on-board the robot's computer, while the interface towards the operator is running on the remote computer.

Fourteen children of preschool age were recruited, six typically developing and eight already diagnosed with ASD. The children were matched by mental age. Gender balance was not considered during participant selection. For all children the session was first contact with the robot, no repeated sessions were considered to avoid learned behaviors. The audio and video recordings of all sessions were analyzed offline and all actions of the robot and social cues exhibited by the

child were extracted and the outcomes are obtained offline by using observed action-observation pairs to calculate the final belief state of the robot. The following task sequence was used for most of the sessions (some sessions were interrupted before finishing all tasks):

1. Play request
2. Functional imitation
3. Joint attention
4. Response to a name call

The play request task is chosen to be performed first to give the child more time to get accustomed to the robot and the examination room as it does not require the child to cooperate with the robot. Then, the objects from the imitation task are used to draw child closer to the robot. After the imitation task the child is expected to be focused on the robot, which enables execution of the joint attention task. The session ends with the response to a name call task with, since the child is likely to be focused on the second robot after the joint attention task.

4 Task MOMDP Models

According to recommendation by AEs, we model a child as being in one of the three states related to low functioning ASD, high functioning ASD and typically developing, that we denote $S^C = \{s_{LA}, s_{HA}, s_{NA}\}$, respectively. Extending our previous work [23], we incorporate set of states used to estimate engagement level of the child S^E with two states (high and low engagement) and set of state used to estimate child verbal activity S^V with also two states (high and low verbal activity). Belief over S^E is updated based on occurrences of child–robot eye contact, while belief over S^V is updated upon the detection of child vocal activity, which does not necessarily coincide with end of task iteration when belief over S^C is updated. In our implementation we maintain




Fig. 2 General structure of a task MOMDP model. The set of all states S is factored into a fully observable set S^{FO} and a partially observable set of states $S^C \times S^E \times S^V$. Robot actions a_k change states in S^T , which is used to track the progress of the task. Social cues that the child may exhibit are coupled in the observation set O^+ . Observations in O^T , which is a subset of O^+ , are used for decision-making within the task

these belief states separately. States S^C , S^E and S^V make up the set of partially observable states of a task model $S^C \times S^E \times S^V$, which is common for all tasks and consists of 12 states. Figure 2 shows a representation of task MOMDP model. With such definition of states, note that the robot actions cannot change the state of the child S^C , which is modeled by using identity state transition matrices for states S^C . As we do not have the data on how the actions of the robot influence the engagement and verbal activity of child, we also set identity transition matrices for S^E and S^V .

Figure 2 also shows that the set of all states S of the task is factored into four subsets, $S^C \times S^E \times S^V \times S^T$, where S^T is the set of observable states used to track the progress of the task. The actions of the robot may change the states in S^T . This interaction of actions and states in S^T is modeled through state transition matrices for each task separately, along with the reward function, in order to achieve the sequence and desired number of repetitions of each action within the task.

In the MOMDP model of a diagnostic task, two sets of observations are distinguished, O^T and O^+ . The observation set O^T consists of observations that directly influence the decision-making in the task (i.e. eye contact for response to a name call, see Table 2) and consequently generate pseudo full observability of the S^T states, as detailed in the following sections. The observation set O^+ contains all observations that are tracked within a given task (including O^T) and is used to build the belief over S^C . Since there are multiple independent observations in the system, belief update for multiple conditionally independent observations is performed by dividing the belief update step into sub-steps such that only one variable is observable in a given sub-step (similar to the approach in [39]). To achieve the proper behavior of the model, the state of the model must not be changed between the sub-steps which translates into requirement that all partial updates to the belief must be performed before pursuing further actions.

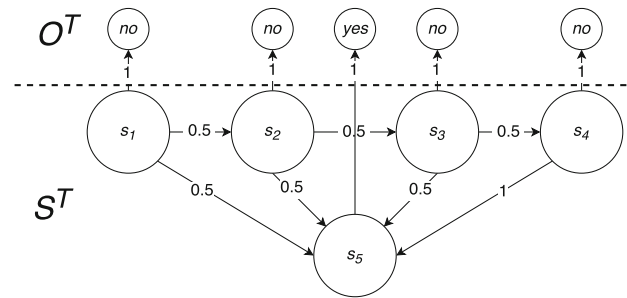


Fig. 3 Set of observable states S^T , transition probabilities and possible observations O^T for actions *call* and *rcall* of the response to a name call task

4.1 Response to a Name Call

This task consists of three calls by name followed by one call that uses a special phrase referring to something dear or interesting to the child. Consequently, the set of actions for this task is $A = \{call, rcall, end\}$, corresponding to a regular call, a call reinforced with special reference and task termination, respectively.

The observable set of states S^T for this task consists of 5 states $s_i \in S^T$, as shown in Fig. 3. The transition functions for each action are defined in matrix form, where each element of the matrix t^a_{ij} defines the following conditional probability of transition:

$$t^a_{ij} = p(s_j \in S^T | s_i \in S^T, a \in A) \tag{1}$$

For the actions *call* and *rcall*, the task is modeled as a variation of left-right-banded type of Markov chain by defining the following transition matrices:

$$\mathbf{T}^{call} = \mathbf{T}^{rcall} = \begin{bmatrix} 0 & 0.5 & 0 & 0 & 0.5 \\ 0 & 0 & 0.5 & 0 & 0.5 \\ 0 & 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \tag{2}$$

The transition function for action *end* is defined as identity matrix, indicating that the action does not change the state of the task.

In the case of the response to a name call, the set of observations O^T is related to eye contact and consists of two observations: $O^T = \{yes, no\}$. The full observability of the S^T is achieved by defining that the observation $o^T = yes$ can only be generated by the state s_5 , while other states can only generate observation $o^T = no$ (see Fig. 3). Such formulation of observation probabilities ensures that the task transitions through states s_1 to s_4 if there is no eye contact, but jumps to s_5 immediately if eye contact is detected.

To achieve the desired task structure, the action *call* is rewarded in states s_1, s_2 and s_3 , the action *rcall* in state s_4 and the action *end* in state s_5 .

The verbal activity of the child is also tracked within the task and the four classes of utterances are grouped into a set O^V . Finally, the set of all observations tracked within the response to a name call task is defined as $O^+ = O^T \times O^V$. In order to build the belief over states S^C , the conditional probabilities of observations in set O^+ need to be specified with respect to the states in S^C and actions A . As there are no studies from which the observation probabilities could be inferred, they are extracted from the experience and expectations of clinicians who are working with children. The formulation of observation probabilities is discussed in Sect. 5.

4.2 Joint Attention

The joint attention task consists of three calls accompanied by turning the head of the robot, one call accompanied by pointing with hand and one attempt of attracting the attention of the child with the other robot. The set of actions for this task is $A = \{turn, point, attract, end\}$. To account for an extra iteration compared to the response to a name call task, six states are used in S^T , with the structure of the observable part of the model shown in Fig. 4.

Again, a left-right banded Markov chain is used as a template, resulting in the following transition probabilities matrices for states in S^T :

$$\mathbf{T}^{turn} = \mathbf{T}^{point} = \mathbf{T}^{attract} = \begin{bmatrix} 0 & 0.5 & 0 & 0 & 0 & 0.5 \\ 0 & 0 & 0.5 & 0 & 0 & 0.5 \\ 0 & 0 & 0 & 0.5 & 0 & 0.5 \\ 0 & 0 & 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \quad (3)$$

The action *turn* is rewarded in states s_1, s_2 and s_3 , the action *point* in state s_4 , the action *attract* in state s_5 and the action *end* in state s_6 . The observation set for the joint attention task is equal to the set used for the response to a




Fig. 4 Set of observable states S^T , transition probabilities and possible observations O^T for actions *turn*, *point* and *attract* of the joint attention task

name call task, except for the fact that eye contact occurrence is detected by the second robot towards which the first one tries to transfer the attention of the child.

4.3 Play Request: Simultaneous Multi-channel Communication Assessment

This is the simplest task to model, as it consists of only one action that is performed three times, regardless of the behavior of the child. Accordingly, the action set is defined as $A = \{perform, end\}$. The state set S^T has four states, while the task observation set O^T is an empty set as there are no observations that influence the decision-making in the task, as shown in Fig. 5. The transition probabilities for action *perform* are:

$$\mathbf{T}^{perform} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (4)$$

Within the task, the robot tracks eye contact, verbal activity and the gestures of the child that could indicate the request for more play. Therefore, the observation set of the task is $O^+ = O^T \times O^E \times O^V$, where $O^E = \{eye\ contact, no\ eye\ contact\}$ are the observations of eye contact, O^V are observations of child utterances as detailed in Sect. 4.1 and $O^G = \{occurred, not\ occurred\}$ are the observations related to the occurrence of gestures that indicate the child’s desire to continue play.

4.4 Functional Imitation

In this task the robot tracks whether the child successfully performs the imitation of drinking and frog jumping, along with eye contact and verbal activity. This means that the observation set O^T consists of two observations, $O^T = \{yes, no\}$, and the set of all observations is the same as in the play request task, namely $O^+ = O^T \times O^E \times O^V$. Regarding the formulation of states in S^T , seven states are needed to achieve the desired task structure, as shown in Fig. 6, since each demonstration is performed three times at the most.

The action set for the imitation task consists of three actions $A = \{frog, drink, end\}$. The action *frog* is rewarded in states s_1, s_2 and s_3 , the action *drink* in states




Fig. 5 Set of observable states S^T and state transition probabilities for action *perform* of the play request task

s_4 , s_5 and s_6 , and the action *end* in state s_7 . The left-right banded Markov chain model for the functional imitation task is encoded in transition probabilities in the following way:

$$\mathbf{T}^{frog} = \mathbf{T}^{drink} = \begin{bmatrix} 0 & 0.5 & 0 & 0.5 & 0 & 0 & 0 \\ 0 & 0 & 0.5 & 0.5 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.5 & 0 & 0.5 \\ 0 & 0 & 0 & 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \tag{5}$$

As can be seen from Fig. 6, upon the observation of correct imitation of the frog jumping gesture, the robot will be sure that the task is now in state s_4 , and switch to the demonstration of drinking gestures. Similarly, if the child imitates drinking successfully, the task will switch to state s_7 in which the task ends.

Finally, actions and observations for each task are summarized in Table 4. The reward function for each task model is specified through immediate rewards with respect to fully observable states to ensure that the structure of a task is kept. The design parameters to be determined for task models are probabilities of all observations in Table 4 taking into account actions (except action *end*) and states S^C .

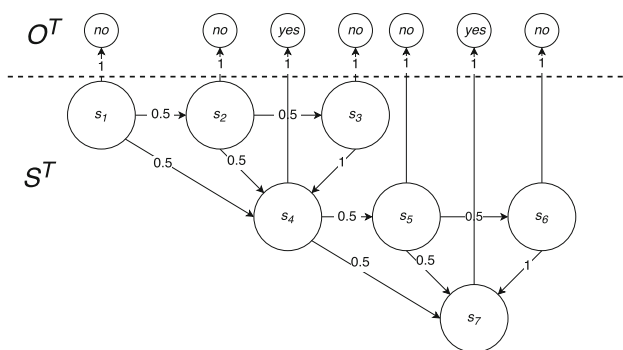


Fig. 6 Set of observable states S^T , transition probabilities and possible observations O^T for actions *turn*, *point* and *attract* of the functional imitation task

Table 4 Task actions and observations

Task	Actions	Observations				
		Eye contact	Joint attention	Verbal activity	Touch	Imitation
RNC	Call, phrase, end	Yes, no	–	No, vocalization, jargon, speech	–	–
JA	Turn, point, attract, end	Yes, no	Yes, no	No, vocalization, jargon, speech	–	–
PR	Play, end	Yes, no	–	No, vocalization, jargon, speech	Yes, no	–
FI	Drink, frog, end	Yes, no	–	No, vocalization, jargon, speech	–	Yes, no

5 Encoding Expert Knowledge in Task Models

In our prior work [23], the observation probabilities were set according to engineer expectations. Herein we improve on those foundations and encode the knowledge of ASD experts in the observation probabilities. For each action within all tasks of the protocol, Gaussian probability density function (*pdf*) estimates are fitted to the histogram of answers by clinical experts from the survey. If in-team experts deem that the observation probability for both low-functioning and high-functioning ASD state are the same for a given action, then unimodal Gaussian *pdf* is used, and the observation value is set to the mode of the *pdf* (i.e. the mean value of expert answers). If in-team experts deem that the observation probability for low-functioning and high-functioning ASD should be different, then the bimodal *pdf* is used, and two modes of such distribution are used as values for observation probability calculation. In such case, the AE team members select which mode corresponds to which state.

5.1 Observation Probabilities for the Response to a Name Call Task

Since the verbal activity of the child is assumed to be the same in all tasks, this section focuses on the determination of eye contact probabilities for two questions in the task from answers to the following four questions:

- Given 10 typically developing children, how many of them respond to the call by name?
- Given 10 typically developing children, how many of them respond to the call using a reference to an object of their liking?
- Given 10 children with ASD, how many of them respond to the call by name?
- Given 10 children with ASD, how many of them respond to the call using a reference to an object of their liking?

We summarize estimates of eye contact observation probabilities for each state of S^C in Table 5.

Table 5 Estimates of eye contact probabilities for actions within response to a name call for each component of state S^C

Action	$p(ec s_{LFA})$	$p(ec s_{HFFA})$	$p(ec s_{TYP})$
<i>call</i>	0.1400 ± 0.049	0.3667 ± 0.094	0.8000 ± 0.254
<i>rcall</i>	0.2667 ± 0.047	0.5200 ± 0.040	0.9125 ± 0.078

Finally, observation matrices for each action in the task are formulated as follows:

$$O^{call} = \begin{matrix} & \begin{matrix} no\ eye\ contact & eye\ contact \end{matrix} \\ \begin{matrix} s_{TYP} \\ s_{HFFA} \\ s_{LFA} \end{matrix} & \begin{bmatrix} 0.1375 & 0.8625 \\ 0.6334 & 0.3666 \\ 0.8600 & 0.1400 \end{bmatrix} \end{matrix}, \quad (6)$$

$$O^{rcall} = \begin{matrix} & \begin{matrix} no\ eye\ contact & eye\ contact \end{matrix} \\ \begin{matrix} s_{TYP} \\ s_{HFFA} \\ s_{LFA} \end{matrix} & \begin{bmatrix} 0.0875 & 0.9125 \\ 0.4800 & 0.5200 \\ 0.7334 & 0.2666 \end{bmatrix} \end{matrix}. \quad (7)$$

5.2 Observation Probabilities for Joint Attention Task

The following six questions are used to extract the observation probabilities for the joint attention task:

- Given 10 typically developing children, how many of them respond to the joint attention request using just speech instructions and head turning?
- Given 10 typically developing children, how many of them respond to the joint attention request using speech instructions, head turning and pointing towards object of interest?
- Given 10 typically developing children, how many of them respond to the object of interest trying to attract their attention?
- Given 10 children with ASD, how many of them respond to the joint attention request using just speech instructions and head turning?
- Given 10 children with ASD, how many of them respond to the joint attention request using speech instructions, head turning and pointing towards object of interest?
- Given 10 children with ASD, how many of them respond to the object of interest trying to attract their attention?

For all actions in this task, the bimodal *pdf* is selected with respect to the probabilities for children with ASD, indicating that there is an expected difference in reactions from low-functioning and high-functioning children with ASD. The estimates of the joint attention observation probabilities for each state of S^C are summarized in Table 6.

Table 6 Estimates of joint attention probabilities for actions of joint attention task for each component of state S^C

Action	$p(ja s_{LFA})$	$p(ja s_{HFFA})$	$p(ja s_{TYP})$
<i>turn</i>	0.0250 ± 0.043	0.2500 ± 0.050	0.8000 ± 0.254
<i>point</i>	0.1800 ± 0.075	0.4667 ± 0.047	0.9375 ± 0.099
<i>attract</i>	0.5200 ± 0.075	0.8333 ± 0.047	0.9625 ± 0.048

Finally, the observation matrices for each action in the joint attention task are formulated as follows:

$$O^{turn} = \begin{matrix} & \begin{matrix} no\ joint\ att. & joint\ attention \end{matrix} \\ \begin{matrix} s_{TYP} \\ s_{HFFA} \\ s_{LFA} \end{matrix} & \begin{bmatrix} 0.200 & 0.800 \\ 0.750 & 0.250 \\ 0.975 & 0.025 \end{bmatrix} \end{matrix}, \quad (8)$$

$$O^{point} = \begin{matrix} & \begin{matrix} no\ joint\ att. & joint\ attention \end{matrix} \\ \begin{matrix} s_{TYP} \\ s_{HFFA} \\ s_{LFA} \end{matrix} & \begin{bmatrix} 0.0625 & 0.9375 \\ 0.5334 & 0.4666 \\ 0.8220 & 0.1780 \end{bmatrix} \end{matrix}, \quad (9)$$

$$O^{attract} = \begin{matrix} & \begin{matrix} no\ joint\ att. & joint\ attention \end{matrix} \\ \begin{matrix} s_{TYP} \\ s_{HFFA} \\ s_{LFA} \end{matrix} & \begin{bmatrix} 0.0360 & 0.9624 \\ 0.1667 & 0.8333 \\ 0.4800 & 0.5200 \end{bmatrix} \end{matrix}. \quad (10)$$

5.3 Observation Probabilities for Play Request

In the play request task, the robot observes verbal activity, tracks eye contact and detects whether the child is performing any action that may suggest a request for more play from the robot. For eye contact in this task the same probabilities determined for the action *attract* in the joint attention task are used, as the robot performs similar acts. For the request part, the experts provided answers to the following questions:

- Given 10 typically developing children, how many of them request more play during one iteration of play request task?
- Given 10 children with ASD, how many of them request more play during one iteration of play request task?

For this particular question, in-team AEs deemed there is no difference in reactions between high and low functioning children with ASD. The mean values extracted from answers to the survey are shown in Table 7.

Table 7 Estimate of probability of child requesting more play for each component of state S^C

Action	$p(req _{SLFA})$	$p(req _{SHFA})$	$p(req _{STYP})$
<i>perform</i>	0.5875 ± 0.154	0.5875 ± 0.154	0.6500 ± 0.335

Observation matrix for the *perform* action is defined as follows:

$$O^{perform} = \begin{matrix} & \begin{matrix} no\ request & request \end{matrix} \\ \begin{matrix} STYP \\ SHFA \\ SLFA \end{matrix} & \begin{bmatrix} 0.3500 & 0.6500 \\ 0.4215 & 0.5785 \\ 0.4215 & 0.5785 \end{bmatrix} \end{matrix}. \quad (11)$$

5.4 Observation Probabilities for Functional Imitation

To estimate the probabilities of a child correctly imitating the demonstrated gesture the following questions were posed:

- Given 10 typically developing children, how many of them correctly imitate demonstrated gesture?
- Given 10 children with ASD, how many of them correctly imitate demonstrated gesture?

As can be inferred from the aforementioned questions, there is no difference in the observation probabilities between the two gestures used in the functional imitation task. The difference in reactions from low-functioning and high-functioning children with ASD are expected in the imitation task and the estimates of imitation observation probabilities for each state of S^C are shown in Table 8.

The observation matrices for two actions in the imitation task are the same and attain the following values:

$$O^{frog} = O^{drink} = \begin{matrix} & \begin{matrix} no\ imitation & imitation \end{matrix} \\ \begin{matrix} STYP \\ SHFA \\ SLFA \end{matrix} & \begin{bmatrix} 0.0875 & 0.9125 \\ 0.6667 & 0.3333 \\ 0.8400 & 0.1600 \end{bmatrix} \end{matrix}. \quad (12)$$

In addition to the gesture imitation, within this task the robot tracks eye contact with the robot, so the probabilities obtained for the action *point* of the joint attention task (see

Table 8 Estimates of imitation probabilities for each component of state S^C

Action	$p(im _{SLFA})$	$p(im _{SHFA})$	$p(im _{STYP})$
<i>frog</i>	0.1600 ± 0.049	0.3333 ± 0.047	0.9125 ± 0.060
<i>drink</i>	0.1600 ± 0.049	0.3333 ± 0.047	0.9125 ± 0.060

Table 6) are used, as the action of demonstrating the gesture and prompting the child to imitate using speech is similar to the pointing gesture in the joint attention task and is expected to cause similar response rate.

5.5 Observation Probabilities for Child’s Verbal Behavior

As already stated, the verbal behavior of a child is modeled as not being dependent on the task or actions, therefore the following questions were used to determine the probability of occurrence of each of four classes of verbal behavior (no verbal activity, vocalizations, jargon and speech):

- Given 10 typically developing children, how many of them do not speak or vocalize during any given interaction?
- Given 10 typically developing children, how many of them vocalize during any given interaction?
- Given 10 typically developing children, how many of them use jargon during any given interaction?
- Given 10 typically developing children, how many of them speak during any given interaction?
- Given 10 children with ASD, how many of them do not speak or vocalize during any given interaction?
- Given 10 children with ASD, how many of them vocalize during any given interaction?
- Given 10 children with ASD, how many of them use jargon during any given interaction?
- Given 10 children with ASD, how many of them speak during any given interaction?

The mean and standard deviation extracted from expert answers are presented in Table 9.

Once the probabilities for verbal activities are obtained, they need to be normalized (sum of probabilities for each state needs to be equal to one). This is necessary for POMDPs (and consequently MOMDPs) as the solver cannot handle simultaneous observations from the same variable (i.e. child exhibiting two classes of verbal behavior within one iteration of action-reaction-observation sequence). The fact that the sum of the probabilities is not equal to one to begin with,

Table 9 Vocal observation probabilities within tasks of the robot-assisted ASD diagnostic protocol for each state in S^C , obtained by surveying experienced ASD clinicians

Observation	$p(obs _{SLFA})$	$p(obs _{SHFA})$	$p(obs _{STYP})$
No activity	0.760 ± 0.120	0.575 ± 0.268	0.088 ± 0.078
Vocalization	0.425 ± 0.156	0.200 ± 0.100	0.513 ± 0.329
Jargon	0.800 ± 0.000	0.438 ± 0.157	0.338 ± 0.303
Speech	0.388 ± 0.145	0.600 ± 0.100	0.888 ± 0.105

is due to humans easily conceiving a situation in which the child produces both speech and vocalization within some interaction, while the POMDP framework considers these observations as mutually exclusive. This cannot be performed by simply scaling probability vector for each state with the inverse of sum of its components, as the scaling factor is not guaranteed to be the same for each vector, which skews the ratio between the probabilities across states which affects the belief update step (i.e. it would change the knowledge representation).

To maintain this ratio of probabilities across the states, the sum of components of the probability vector for each state is calculated. For vocal observations in Table 9, all values in each column are added. Then, all entries in the matrix are scaled with the maximum of obtained column sums. A residual observation is introduced into the model and is used to collect the remaining probability for each state, resulting in the final observation probabilities matrix that is used in the MOMDP model:

$$O^{ver.} = \begin{matrix} & \begin{matrix} no\ act. & voc. & jarg. & speech & res. \end{matrix} \\ \begin{matrix} s_{TYP} \\ s_{HFA} \\ s_{LFA} \end{matrix} & \begin{bmatrix} 0.037 & 0.216 & 0.142 & 0.374 & 0.231 \\ 0.242 & 0.084 & 0.184 & 0.253 & 0.236 \\ 0.320 & 0.179 & 0.337 & 0.163 & 0.0001 \end{bmatrix} \end{matrix} \quad (13)$$

As all other observation sets are binary, this procedure is not necessary to formulate consistent matrices.

5.6 Updating Belief with Respect to Engagement and Verbal Activity

Through the survey, the experts also provided estimates on how many eye contacts are needed to deem the child to be engaged with the examiner, and the answer was three. A similar question was posed regarding verbal activity and the experts estimate that the number of verbal actions needed to deem the child verbally active was five. As the eye contact detection on the robot runs somewhat faster than sound classification, five is used as the target number of occurrences for both eye contact and speech of the child to deem the child engaged with the robot and verbally active. In terms of the belief state of the robot, targets of $b(s_{HE}) > 0.9$ and $b(s_{HV}) > 0.9$ are set to be reached after five detections.

At this stage, there is no data on how the actions of the robot affect the engagement and verbal activity of the child, so identity state transition probability matrices are used indicating that actions of the robot cannot change states in S^E and S^V . Ideally, the initial belief over s_{HE} and s_{HV} should be zero, but that coupled with identity state transition matrices would result in belief state not changing with new observations. Therefore, the initial belief state is set to the following values:

$$b(s_{HE})^0 = b(s_{HV})^0 = 0.1 \quad (14)$$

Now, the observation probabilities remain to be determined. With the goal of having $b(s_{HE}) > 0.9$ after five detections of eye contact, the following eye contact detection probabilities are set:

$$O^E = \begin{matrix} & \begin{matrix} no\ eye\ contact & eye\ contact \end{matrix} \\ \begin{matrix} s_{HE} \\ s_{LE} \end{matrix} & \begin{bmatrix} 0.2 & 0.8 \\ 0.8 & 0.2 \end{bmatrix} \end{matrix} \quad (15)$$

which result in $b(s_{HE}) = 0.9118$ after five detections of eye contact. To achieve similar belief values for verbal activity after five detections of speech and $b(s_{HE}) \approx 0.5$ after five detections of jargon, the following verbal observation probabilities are used:

$$O^V = \begin{matrix} & \begin{matrix} no\ act. & voc. & jargon & speech \end{matrix} \\ \begin{matrix} s_{HV} \\ s_{LV} \end{matrix} & \begin{bmatrix} 0.10 & 0.20 & 0.30 & 0.40 \\ 0.48 & 0.30 & 0.12 & 0.10 \end{bmatrix} \end{matrix} \quad (16)$$

Updating the belief for five occurrences of speech results in $b(s_{HV}) = 0.9118$, which is the same value as for $b(s_{HE})$. This is expected as the ratio of observation probabilities between states for speech and eye contact are the same. The formulation of observation probabilities in (16) indicates that the evolution of belief over the child verbal activity results in different estimates of verbal activity for different observations detected. It can also be observed that only detections of jargon and speech contribute to higher verbal activity estimates, while detection of vocalization lowers the belief over s_{HV} as vocalizations are generally considered to be *pre-verbal* form of communication.

With the expert knowledge encoded in the transition and observation probabilities of the task MOMDP models, the remaining question to be answered is whether sequential updates of the robot belief using multiple observations generate outcomes that have diagnostic validity and to which degree can robot’s belief after the end of a task be used as a measure of child’s behavior during task administration.

6 Monte Carlo Simulation of Child–Robot Interaction

To facilitate the simulation of the interaction, a stochastic behavioral model is formulated in which any action of the robot *samples* the response of the child according to the probability distribution of social cues obtained from experts in the survey. Due to the lack of data, the engagement and verbal level estimation are omitted from investigations and consequently from child models and simulation is used to evaluate diagnostic validity of tasks and to tune the amount of

information each observation brings into the robot's belief. Three stochastic child models are needed, one for each of the considered child types. Each model consists of pairs of actions and vectors of observation probabilities from all task MOMDP models. The simulation of a task is performed by repeating the following steps N times:

1. Set initial belief state b to uniform over states S^C .
2. Select next action a using b and policy of the task.
3. If the next action is *end*, the task is finished. Save b for further analysis and go to step 1 if number of task instances simulated is smaller than N . If task has been simulated N times, end the simulation.
4. For action a select observation probability vector for social cues considered in the task and perform *roulette wheel selection* to obtain the observations o for this step
5. Update belief using b , a and o and go to step 2.

Step 4 of one simulation iteration contains a *roulette wheel selection* of the observation which will be detected. The roulette wheel selection, also known as the *Fitness proportionate selection*, is a common operator in genetic algorithms. It is used to select chromosomes of the genetic algorithm that are selected for recombination based on the fitness of each chromosome. The more fit the chromosome is to some fitness function, the more likely it is to be selected. For child behavioral models, chromosomes are replaced with observations and fitness is replaced with observation probabilities in order to simulate child–robot interaction.

In the most common implementation of the roulette wheel selection, which is also adopted in this work since there are not many observations, the first step is to formulate a cumulative distribution function (CDF) over the list of observations using observation probabilities. This operation is equivalent to setting the number of bins on a roulette wheel to the number of observations and their widths according to their respective probabilities. Next, a uniform random number n in the range $[0, 1)$ is generated using a random number generator. Finally, taking the inverse of CDF for n gives out the observation that the model generates.

6.1 Scaling the Amount of Information Each Observation Brings into the Belief of the Model

Simulating the child–robot interaction showed that in some cases the outcome of the task defined by the belief of the robot over states of the child can be decided solely by the last action–observation pair and not by the sequence of actions and observations during the whole task. The possibility of abrupt change of belief in only one iteration indicates that the amount of information that the observations bring into the task model needs to be scaled down. This can be achieved by flattening the probabilities of observations within an observa-

tion set across all states. This results in belief being updated by smaller amount, preventing abrupt changes. If the amount of information each observation brings into the task model is denoted $p_s \in (0.5, 1.0)$,¹ the observation probability $p(o|s, a)$ can be flattened in the following way:

$$p'(o|s, a) = p_s \cdot p(o|s, a) + (1 - p_s) \cdot [1 - p(o|s, a)]. \quad (17)$$

Different amounts of information for each observation can be specified by using different values of p_s for different observation sets, such as observations that directly influence the decision-making within tasks (O^T) and additional observations in the task ($O^A = O^+ \setminus O^T$). The amount of information of the observations in O^T is denoted p_s^T , and p_s^A denotes the importance of the additional observations O^A . After several iterations of tuning, the amounts of information in observation sets is set to the following values:

$$p_s^T = 0.9, \quad (18)$$

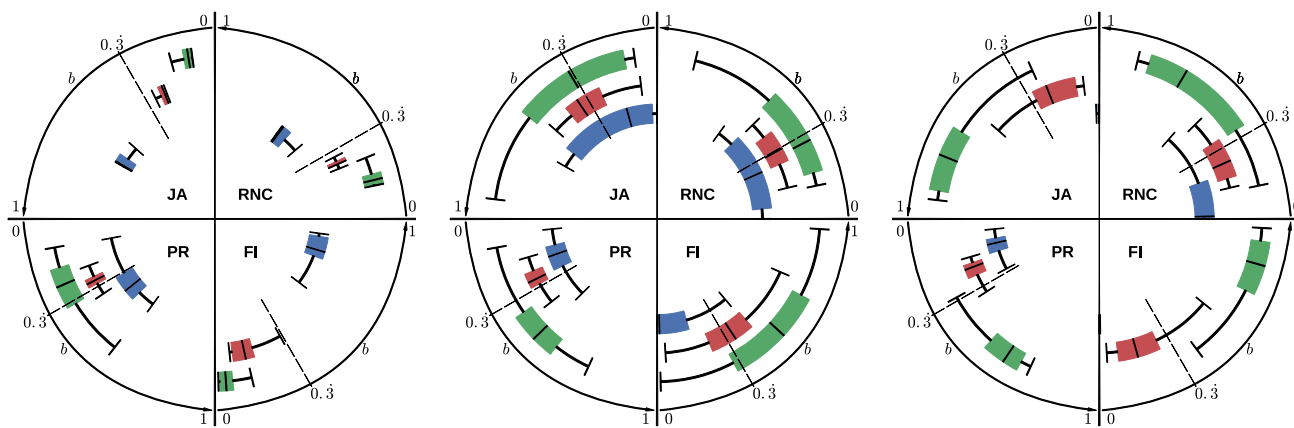
$$p_s^A = 0.7. \quad (19)$$

6.2 Evaluation of Diagnostic Capabilities of Task Models

The distributions of outcomes after simulating 10000 iterations of all tasks for all the child behavioral models are summarized in Fig. 7.

Figure 7a shows that three of the tasks are capable of identifying typical behavior, albeit with lower confidence than identifying low-functioning behavior, if compared to Fig. 7c. This is to be expected since typical behavior results in lower number of iterations in the task (task ends when child responds), so there are fewer opportunities to update the belief over the state of the child. The outcomes of play request task from Fig. 7a, while showing trend towards identifying typical behavior should be classified as inconclusive as belief is nearly an uniform distribution over considered states. Outcomes for the simulation of high-functioning ASD behavior, shown in Fig. 7b, indicate that task models are not suited for identifying such behavior, as outcomes are either inconclusive or classify behavior to be more similar to that expected from a child with low-functioning ASD. To measure how clearly the belief of the robot can identify an underlying type of behavior, the Hellinger distance [40] is used to measure the distance between distributions of $b(STYP)$, $b(SHFA)$ and $b(SLFA)$. The Hellinger distance is used to quantify how

¹ Value $p_s = 0.5$ results in uniform distribution of observation probabilities while value $p_s = 1.0$ leaves probabilities unchanged. Values $p_s > 1$ increase the difference in probabilities increasing the amount of information the observation brings into the task. Values $p_s < 0.5$ result in probabilities that update belief in the wrong direction.



(a) Outcomes of diagnostic task simulations with the model of a typically developing child.

(b) Outcomes of diagnostic task simulations with the model of a child with high-functioning ASD.

(c) Outcomes of diagnostic task simulations with the model of a child with low-functioning ASD.

Fig. 7 Distributions of the model’s belief at the end of each task. Each quadrant in a graph represents the outcomes of one of the tasks (tasks are labeled according to Sect. 2). Box plots in the inner-most circle (blue color) represent distributions of the model’s belief that the child exhibited behavior similar to that expected from a typically developing child, $b(s_{TYP})$. Outer-most circle box plots (green color) show distributions of the model’s belief that the child exhibited behavior similar to that

expected from a child with low-functioning ASD, $b(s_{LFA})$. The box plots in the middle band (red color) show distributions of the model’s belief that the child exhibited behavior similar to that expected from a child with high-functioning ASD, $b(s_{HFA})$. Dashed lines mark a point in the belief space of the model where all components of belief over states in S^C are the same (i.e. the outcome of a task is inconclusive if all components of belief state are near this line). (Color figure online)

similar are two probability distributions, P and Q . Distance can attain values in range $[0, 1]$, with maximum distance 1 describing a scenario in which P assigns probability zero to every outcome to which Q assigns some probability, and vice-versa. For two discrete probability distributions, or in this case two normalized histograms with the same amount of bins, $P = (p_1, \dots, p_n)$ and $Q = (q_1, \dots, q_n)$, the Hellinger distance is defined as follows [41]:

$$H(P, Q) = \frac{1}{\sqrt{2}} \sqrt{\sum_{i=1}^n (\sqrt{p_i} - \sqrt{q_i})^2}. \tag{20}$$

The Hellinger distances between the task outcomes for each of the child models are shown in Table 10. As already mentioned, the higher values correspond to the greater dis-

tance between histograms meaning less overlap between outcomes of the task.

If the child exhibits the typical behavior, there is practically no overlap of $b(s_{TYP})$ with other two components of the belief state for response to name call and joint attention tasks. The distances between belief components for imitation task attain slightly lower values but also show that there is no significant overlap of outcomes, indicating that the imitation task is also suitable for the detection of the typical behavior of the child. On the other hand, the distances between the outcomes of the play request task show a significant overlap, which confirms that the task is not suitable for detection of the typical behavior.

The outcomes of simulation with a model of a child with high functioning ASD show the most overlap, which is obvious from Fig. 7 but is also indicated by lower values of distances in Table 10. It can be concluded that none of the

Table 10 Hellinger distance between distributions of task outcomes for $p_s^T = 0.9$ and $p_s^A = 0.7$

Model	H	Task			
		Name call	Joint attention	Play request	Imitation
TYP	$H(b(s_{TYP}), b(s_{HFA}))$	0.952	0.923	0.445	0.873
	$H(b(s_{TYP}), b(s_{LFA}))$	0.953	0.954	0.609	0.923
HFA	$H(b(s_{TYP}), b(s_{HFA}))$	0.785	0.741	0.287	0.695
	$H(b(s_{LFA}), b(s_{HFA}))$	0.707	0.458	0.425	0.209
LFA	$H(b(s_{TYP}), b(s_{LFA}))$	0.781	0.855	0.852	0.891
	$H(b(s_{LFA}), b(s_{HFA}))$	0.708	0.798	0.845	0.651

tasks are particularly well suited for detection of highly functional autistic behavior. If the child exhibits autistic behavior as simulated using model of a child with low functioning ASD, all tasks are successful in correctly predicting the child type with small overlap between outcomes.

To summarize, the simulation of child–robot interaction and analysis of outcomes confirmed that the task models with observation probabilities extracted from expert knowledge can, for the most part, successfully differentiate between the autistic behavior and the behavior characteristic for a typically developing children but fail to differentiate between the degrees of autism severity, as most of the outcomes for all three models end at the either end of the spectrum. To some degree, this is expected as there is no observation in the model that is characteristic to s_{HFA} . In all observation probability matrices in Sect. 5, there is no observation probability associated to s_{HFA} that is greater than probabilities assigned to other two states in S^C , which means that for any observation detected, the belief of the robot will steer towards s_{LFA} and s_{TYP} more than towards s_{HFA} .

7 Experimental Sessions with Children in Clinical Setting

As already mentioned, fourteen children of preschool age were recruited, six typically developing and eight already diagnosed with ASD. During the sessions, it was observed that typically developing children sometimes exhibit autistic behavior, and vice versa. This was also shown in our previous work on imitation [26], with possible reasons being that all the children in the ASD group are undergoing intervention and are trained to perform similar tasks almost daily and exhibited no anxiety towards the robot while typically developing children were more shy and wary of the robot. To properly validate the proposed task MOMDP models, it is important to analyze the outcomes of the protocol based on the exhibited behavior, not the child diagnosis. In order to do so, the sessions with children were transcribed as a sequence of actions and observations, anonymized and shared with AE team members. AE team members were asked to evaluate the transcript of the child behavior and conclude whether the behavior within each task is more similar to that of a typically developing child or to that of a child with ASD. If not confident in making the assessment, AEs were encouraged to classify task results as inconclusive. Results of AE classifications are summarized in Table 11.

As can be seen from Table 11, not all children performed all actions and all tasks, some due to the various robot failures some due to children themselves being afraid or not wanting to cooperate with the robot. Results of eleven sessions of the response to a name call task are shown in Fig. 8. The graphs show mean values of belief at each iteration of a task as points,

Table 11 Number of sessions classified in each of the behavior classes considered by ASD experts for every task of the protocol

	Typical	ASD	Inconclusive	Total
RNC	6	4	1	11
JA	1	8	3	12
PR	5	3	6	14
FI	4	4	3	11

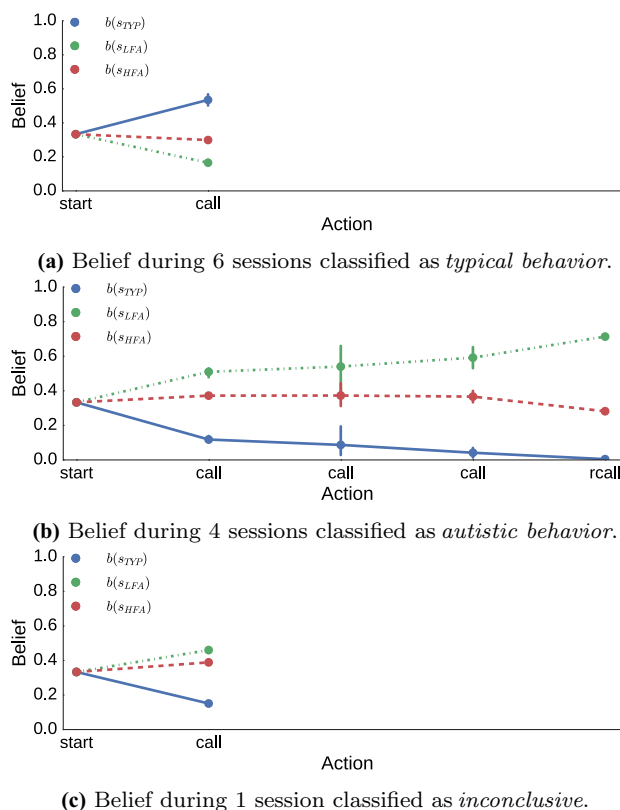


Fig. 8 Belief of the MOMDP task model over child states during the response to a name call task. (Color figure online)

while the vertical bars at those points show the spread of the belief at the same iteration. If there is no bar, it means that the spread of belief is not significant in those points, which occurs if the sequence of actions and observations is similar up to that point. Mean values of belief in subsequent points are connected with a line to better visualize the evolution of belief between iterations.

Evolution of the robot’s belief in Fig. 8a shows that the robot correctly identifies typical behavior. In all sessions the children responded immediately and the only difference was in using speech, which results in small variance of belief at the end of a task. Similar can be observed for autistic behavior, for which the evolution of the robot’s belief is shown in Fig. 8b. Again, the variance in belief during the task is generated by differences in verbal activity. The belief shown

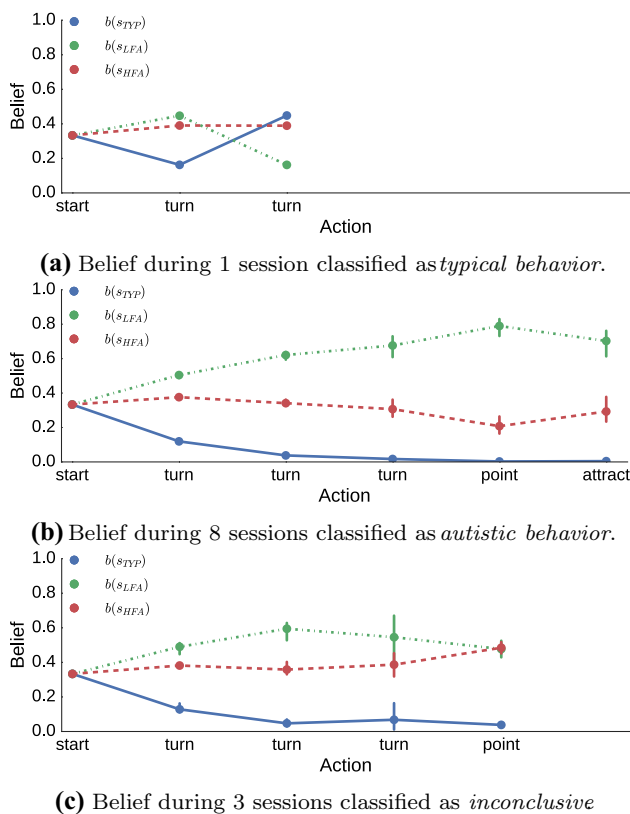


Fig. 9 Belief of the MOMDP task model over child states during the joint attention task. (Color figure online)

in Fig. 8c cannot be used to draw meaningful conclusions as it presents only one session, but suggests that in cases in which the human is not confident enough to make the assessment, the robot may be biased towards classifying behavior as autistic.

Figure 9 shows the belief of the robot during 12 sessions of the joint attention task. Although result of only one session, Fig. 9a provides valuable insight into the evolution of the robot’s belief. In this session, the child did not respond in the first iteration, but responded in the second one and used speech (which is deemed to correspond with typical behavior in the models) so the robot immediately changed its estimate. Graph in Fig. 9b shows that the robot can correctly identify behavior classified by ASD experts as autistic, while Fig. 9c again shows that the sessions classified as inconclusive by humans are classified as autistic behavior by the robot. However, comparing Fig. 9b, c, one can observe the difference in graphs, and that the robot fails to clearly identify the behavior as low-functioning ASD, indicating that the robot is also less confident in the results from these sessions.

The robot’s belief during the play request task is shown in Fig. 10, with graphs confirming that this task is the least informative, as suggested by simulation results from Fig. 7. While graphs in Fig. 10a, b confirm that the robot can identify

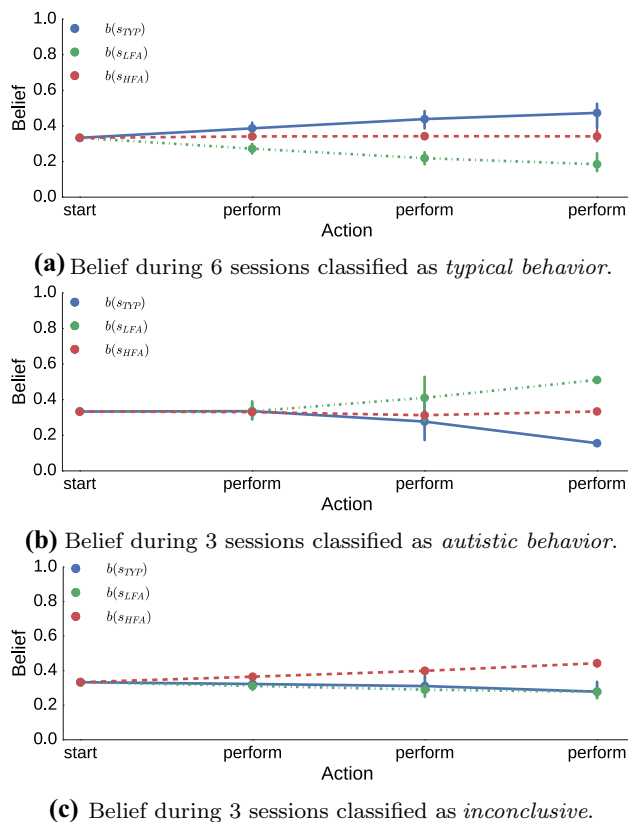


Fig. 10 Belief of the MOMDP task model over child states during the play request task. (Color figure online)

the behavior of the child in accordance with assessment of a human, although with not much confidence, the graph in Fig. 10c shows that the robot tends to identify the behaviors deemed to be inconclusive by humans as similar to those expected from a child with high-functioning ASD. Similar pattern can be observed in Fig. 11. Again, both Fig. 11a, b confirm that the robot can correctly identify the behavior of the child in the same way a human does, while Fig. 11c shows that for the sessions deemed to be inconclusive the robot infers that the behavior is related to high-functioning ASD.

More data is required to draw conclusions in cases in which the behavior of children is ambiguous. Although there is not enough data to claim that sessions deemed inconclusive by ASD experts reflect high-functioning ASD behavior, it is interesting to consider that in those cases the results from Figs. 10c and 11c suggest that, even though there are no observations in the model that are characteristic to high-functioning autistic behavior, the models are capable to infer that the behavior may be similar to that expected from a child with high-functioning ASD.

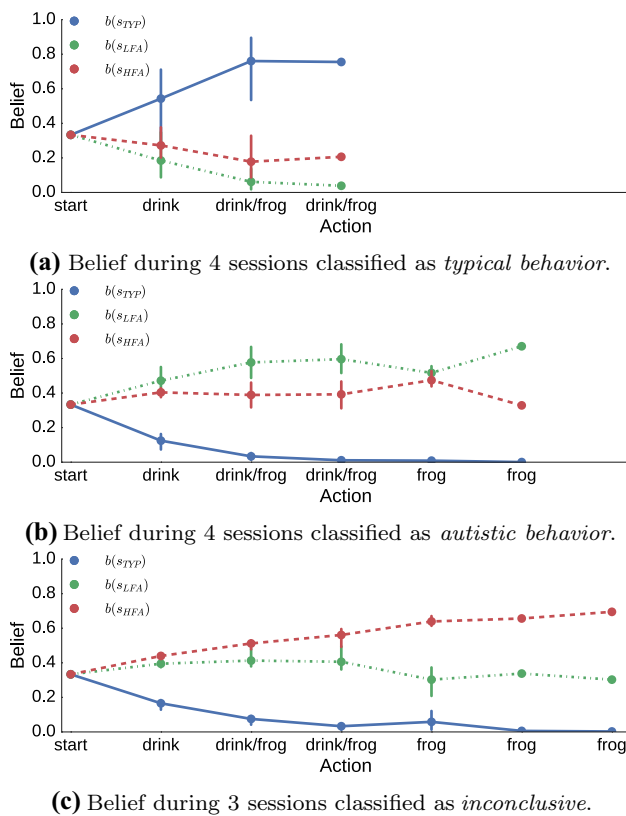


Fig. 11 Belief of the MOMDP task model over child states during the functional imitation task. (Color figure online)

8 Conclusion

In this paper, we presented a novel method to design MOMDP task models for controlling and evaluating the child–robot interaction within the robot-assisted ASD diagnostic protocol. Each task of the protocol is modeled as an MOMDP. Observation probabilities of each MOMDP model are set according to answers gathered by surveying ASD experts. Same survey is used to formulate stochastic models that are used to simulate the interaction. Simulation results validated task models as capable of identifying the underlying behavior of the child but also enabled fine-tuning of some parameters in task models. Finally, the task models were validated through experimental session with six typically developing children and eight children with ASD. The MOMDP model’s belief at the end of each task was compared to assessment of anonymized transcript of the interaction by human experts. The comparison showed that the model’s belief can be used as an automatic evaluator of the child’s behavior. These results enable development of a protocol model which will enable the robot to adapt to different behavior of a child and choose sequence of tasks in a given session to maximize information gathered from the interaction.

The major drawback of the MOMDP task models as formulated in this work is the dependency on the accuracy of social cue detection, which was not high enough to perform the experimental evaluation in a way in which the robot is fully autonomous. Rather, the human operator provided the robot with corrections of observations of social cues in cases where automatic detection failed. With the reliance on social cues and capabilities of NAO robots in mind, there is little probability that the robots will be performing the tasks autonomously in the near future. A more likely scenario is the robots being a part of a smart room for autism diagnosis in which multiple sensors will be installed to aid the robot in tracking the child and observing its behavior. Since the proposed methodology of formulating a task model can be used with any task from the ADOS, there is a possibility of extending the protocol with new tasks, for which the most prominent candidate is the task of free play in which the robot just observes a child playing with toys and encodes the preferences of the child. A more immediate future work is likely to include more sessions with children with emphasis on evaluating whether the task models can indeed identify high-functioning behavior.

Acknowledgements The authors would like to thank the children and their parents for volunteering in experimental sessions and Damjan, Maja, Jasmina, Sanja, Petra, Dina and Omar for their immeasurable contribution.

Funding This work has been fully supported by the Croatian Science Foundation through the ADORE project (HRZZ-93743-2014).

Compliance with Ethical Standards

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

References

- Baio J, Wiggins L, Christensen DL, Maenner MJ, Daniels J, Warren Z, Kurzius-Spencer M, Zahorodny W, Rosenberg CR, White T, Durkin MS, Imm P, Nikolaou L, Yeargin-Allsopp M, Lee L-C, Harrington R, Lopez M, Fitzgerald RT, Hewitt A, Pettygrove S, Constantino JN, Vehorn A, Shenouda J, Hall-Lande J, Van Naarden Braun K, Dowling NF (2018) Prevalence of autism spectrum disorder among children aged 8 years-autism and developmental disabilities monitoring network, 11 sites, United States, 2014. *MMWR. Surveill Summ* 67(6):1–23. <https://doi.org/10.15585/mmwr.ss6706a1>
- Roelfsema MT, Hoekstra RA, Allison C, Wheelwright S, Brayne C, Matthews FE, Baron-Cohen S (2011) Are autism spectrum conditions more prevalent in an information-technology region? A school-based study of three regions in the Netherlands. *J Autism Dev Disord* 42(5):734–739. <https://doi.org/10.1007/s10803-011-1302-1>
- Kim YS, Leventhal BL, Koh YJ, Fombonne E, Laska E, Lim EC, Cheon KA, Kim SJ, Kim YK, Lee H, Song DH, Grinker RR (2011) Prevalence of autism spectrum disorders in a total popu-

- lation sample. *Am J Psychiatry* 168(9):904–912. <https://doi.org/10.1176/appi.ajp.2011.10101532>
4. Klin A, Lang J, Chicchetti V, Volkmar F (2000) Interrater reliability of clinical diagnosis and DSM-IV criteria for autistic disorder: results of the DSM-IV autism field trial. *J Autism Dev Disord* 30(2):163–167
 5. Crane L, Chester JW, Goddard L, Henry LA, Hill E (2015) Experiences of autism diagnosis: a survey of over 1000 parents in the united kingdom. *Autism* 20(2):153–162. <https://doi.org/10.1177/1362361315573636>
 6. Lavelle TA, Weinstein MC, Newhouse JP, Munir K, Kuhlthau KA, Prosser LA (2014) Economic burden of childhood autism spectrum disorders. *Pediatrics* 133(3):e520–e529. <https://doi.org/10.1542/peds.2013-0763>
 7. Daniels AM, Como A, Hergüner S, Kostadinova K, Stosic J, Shih A (2017) Autism in southeast Europe: a survey of caregivers of children with autism spectrum disorders. *J Autism Dev Disord* 47(8):2314–2325. <https://doi.org/10.1007/s10803-017-3145-x>
 8. Scassellati B (2005) Quantitative metrics of social response for autism diagnosis. In: *IEEE international workshop on robot and human interactive communication, 2005 (ROMAN 2005)*, pp 585–590
 9. Aresti-Bartolome N, Garcia-Zapirain B (2014) Technologies as support tools for persons with autistic spectrum disorder: a systematic review. *Int J Environ Res Public Health* 11(8):7767–7802. <https://doi.org/10.3390/ijerph110807767>
 10. Pennisi P, Tonacci A, Tartarisco G, Billeci L, Ruta L, Gangemi S, Pioggia G (2015) Autism and social robotics: a systematic review. *Autism Res* 9(2):165–183. <https://doi.org/10.1002/aur.1527>
 11. Petric F, Hrvatinić K, Babic A, Malovan L, Miklic D, Kovacic Z, Cepanec M, Stosic J, Simlesa S (2014) Four tasks of a robot-assisted autism spectrum disorder diagnostic protocol: first clinical tests. In: *2014 IEEE global humanitarian technology conference (GHTC)*, pp 510–517. <https://doi.org/10.1109/GHTC.2014.6970331>
 12. Boccanfuso L, Scarborough S, Abramson RK, Hall AV, Wright HH, O’Kane JM (2016) A low-cost socially assistive robot and robot-assisted intervention for children with autism spectrum disorder: field trials and lessons learned. *Auton Robots*. <https://doi.org/10.1007/s10514-016-9554-4>
 13. Huijnen CAGJ, Lexis MAS, de Witte LP (2016) Matching robot KASPAR to autism spectrum disorder (ASD) therapy and educational goals. *Int J Soc Robot* 8(4):445–455. <https://doi.org/10.1007/s12369-016-0369-4>
 14. Taheri A, Meghdari A, Alemi M, Pouretamad H (2017) Human–robot interaction in autism treatment: a case study on three pairs of autistic children as twins, siblings, and classmates. *Int J Soc Robot* 10(1):93–113. <https://doi.org/10.1007/s12369-017-0433-8>
 15. David DO, Costescu CA, Matu S, Szentagotai A, Dobrean A (2018) Developing joint attention for children with autism in robot-enhanced therapy. *Int J Soc Robot* 10(5):595–605. <https://doi.org/10.1007/s12369-017-0457-0>
 16. Ismail LI, Verhoeven T, Dambre J, Wyffels F (2018) Leveraging robotics research for children with autism: a review. *Int J Soc Robot*. <https://doi.org/10.1007/s12369-018-0508-1>
 17. d Souza PEU, Chanel CPC, Dehais F (2015) Momdp-based target search mission taking into account the human operator’s cognitive state. In: *2015 IEEE 27th international conference on tools with artificial intelligence (ICTAI)*, pp 729–736. <https://doi.org/10.1109/ICTAI.2015.109>
 18. Nikolaidis S, Zhu YX, Hsu D, Srinivasa S (2017) Human–robot mutual adaptation in shared autonomy. In: *Proceedings of the 2017 ACM/IEEE international conference on human–robot interaction (HRI’17)*. ACM, New York, pp 294–302. <https://doi.org/10.1145/2909824.3020252>
 19. Anderson B, Moore A (2005) Active learning for hidden Markov models: Objective functions and algorithms. In: *Proceedings of the 22nd international conference on machine learning (ICML’05)*. ACM, New York, pp 9–16. <https://doi.org/10.1145/1102351.1102353>
 20. Atrash A, Pineau J (2009) A Bayesian reinforcement learning approach for customizing human–robot interfaces. In: *Proceedings of the 14th international conference on intelligent user interfaces (IUI ’09)*. ACM, New York, pp 355–360. <https://doi.org/10.1145/1502650.1502700>
 21. Doshi-Velez F, Pineau J, Roy N (2012) Reinforcement learning with limited reinforcement: using bayes risk for active learning in POMDPs. *Artif Intell* 187–188:115–132. <https://doi.org/10.1016/j.artint.2012.04.006>
 22. Petric F, Tolić D, Miklić D, Kovačić Z, Cepanec M, Šimleša S (2015) Towards a robot-assisted autism diagnostic protocol: modelling and assessment with POMDP. In: *Intelligent robotics and applications*. Springer, Berlin, pp 82–94. https://doi.org/10.1007/978-3-319-22876-1_8
 23. Petric F, Miklić D, Kovačić Z (2018) POMDP-based coding of child–robot interaction within a robot-assisted ASD diagnostic protocol. *Int J Hum Robot* 15(02):1850011. <https://doi.org/10.1142/s0219843618500111>
 24. Lord C, Rutter M, Goode S, Heemsbergen J, Jordan H, Mawhood L, Schopler E (1989) Autism diagnostic observation schedule: a standardized observation of communicative and social behavior. *J Autism Dev Disord* 19(2):185–212. <https://doi.org/10.1007/bf02211841>
 25. Petric F, Miklić D, Kovačić Z (2016) Probabilistic eye contact detection for the robot-assisted ASD diagnostic protocol. In: *Proceedings of the fifth Croatian computer vision workshop (CCVW 2016)*, pp 3–8
 26. Petric F, Miklic D, Cepanec M, Cvitanovic P, Kovacic Z (2017) Functional imitation task in the context of robot-assisted autism spectrum disorder diagnostics: preliminary investigations. In: *2017 26th IEEE international symposium on robot and human interactive communication (RO-MAN)*. IEEE. <https://doi.org/10.1109/roman.2017.8172498>
 27. Kokot M, Petric F, Cepanec M, Miklić D, Bejić I, Kovačić Z (2018) Classification of child vocal behavior for a robot-assisted autism diagnostic protocol. In: *2018 26th Mediterranean conference on control and automation (MED)*, pp 1–6. <https://doi.org/10.1109/MED.2018.8443030>
 28. Dawson G, Toth K, Abbott R, Osterling J, Munson J, Estes A, Liaw J (2004) Early social attention impairments in autism: social orienting, joint attention, and attention to distress. *Dev Psychol* 40(2):271–283
 29. Jones W, Carr K, Klin A (2008) Absence of preferential looking to the eyes of approaching adults predicts level of social disability in 2-year-old toddlers with autism spectrum disorder. *Arch Gen Psychiatry* 65(8):946–954
 30. Hauck M, Fein D, Waterhouse L, Feinstein C (1995) Social initiations by autistic children to adults and other children. *J Autism Dev Disord* 25(6):579–595. <https://doi.org/10.1007/bf02178189>
 31. Stone WL, Caro-Martinez LM (1990) Naturalistic observations of spontaneous communication in autistic children. *J Autism Dev Disord* 20(4):437–453. <https://doi.org/10.1007/bf02216051>
 32. Carpenter M, Nagell K, Tomasello M, Butterworth G, Moore C (1998) Social cognition, joint attention, and communicative competence from 9 to 15 months of age. *Monographs of the society for research in child development*, p 1-174
 33. Charman T, Swettenham J, Baron-Cohen S, Cox A, Baird G, Drew A (1997) Infants with autism: an investigation of empathy, pretend play, joint attention, and imitation. *Dev Psychol* 33(5):781

34. Ong SCW, Png SW, Hsu D, Lee WS (2010) Planning under uncertainty for robotic tasks with mixed observability. *Int J Robot Res* 29(8):1053–1068. <https://doi.org/10.1177/0278364910369861>
35. Brock O, Trinkle J, Ramos F (2009) SARSOP: efficient point-based POMDP planning by approximating optimally reachable belief spaces. MIT Press, Cambridge, pp 65–72
36. Azenkot S, Feng C, Cakmak M (2016) Enabling building service robots to guide blind people a participatory design approach. In: 2016 11th ACM/IEEE international conference on human-robot interaction (HRI), pp 3–10. <https://doi.org/10.1109/HRI.2016.7451727>
37. Sequeira P, Alves-Oliveira P, Ribeiro T, Di Tullio E, Petisca S, Melo FS, Castellano G, Paiva A (2016) Discovering social interaction strategies for robots from restricted-perception wizard-of-oz studies. In: The eleventh ACM/IEEE international conference on human robot interaction (HRI'16). IEEE Press, Piscataway, pp 197–204. <http://dl.acm.org/citation.cfm?id=2906831.2906866>
38. (2015) Nao Software Documentation. Aldebaran Robotics, v2.1 edn. <https://community.aldebaran-robotics.com/doc>
39. Hoey J, Poupart P (2005) Solving pomdps with continuous or large discrete observation spaces. In: Proceedings of the 19th international joint conference on artificial intelligence (IJCAI'05). Morgan Kaufmann, San Francisco, pp 1332–1338. <http://dl.acm.org/citation.cfm?id=1642293.1642505>
40. Hellinger E (1909) Neue begründung der theorie quadratischer formen von unendlichvielen veränderlichen. *Journal für die reine und angewandte Mathematik (Crelle's Journal)* 1909(136). <https://doi.org/10.1515/crll.1909.136.210>
41. Nikulin M (2001) Hellinger distance—encyclopedia of mathematics. European Mathematical Society/Springer, Berlin

Franjo Petric received his PhD degree in electrical engineering from the Faculty of Electrical Engineering and Computing, University of Zagreb, Zagreb, Croatia in 2018. Since 2012, he has been working as a Research Associate with the Laboratory for Robotics and Intelligent Control Systems, University of Zagreb, taking part in multiple national and international projects (EC-SAFEMOBIL, Assisi|bf, Aerotwin, ADORE). His current research interests include human-robot interaction, artificial intelligence, computer vision and software engineering for robotics with emphasis on child-robot interaction and use of humanoid robots with children with ASD. He has authored or co-authored multiple papers in journals and conferences, along with two book chapters.

Zdenko Kovacic was born in 1958. He is currently a Full Professor with the Faculty of Electrical Engineering and Computing, University of Zagreb, and the Head of the Laboratory for Robotics and Intelligent Control Systems. He has authored or co-authored two scientific monographs, one University textbook, and many journal articles and conference papers. His current research interests include robotics, control and automation with main contributions in robot control, co-operative robotic systems, intelligent control, service robotics, humanoids, and walking robots. He has earned recognition from industry partners for being the Principal Investigator of over 30 successfully completed international and Croatian Research and Development Projects. Prof. Kovacic is a recipient of the University of Zagreb Award Fran Bosnjakovic in 2013 for his distinctive scientific efforts in advancing and developing control, flexible automation, and robotics systems. He was the President of the Croatian Robotics Society from 2005 to 2010, the President of the Croatian IEEE Robotics and Automation Chapter from 2010 to 2013, and Elected President of the Croatian Robotic Association from 2012 to 2015.

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