Towards A Robot-Assisted Autism Diagnostic Protocol: Modelling and Assessment with POMDP

Frano Petric[1(✉)], Domagoj Tolić[1], Damjan Miklić[1], Zdenko Kovačić[1], Maja Cepanec[2], and Sanja Šimleša[2]

[1] Faculty of Electrical Engineering and Computing, University of Zagreb, Zagreb, Croatia
{frano.petric,domagoj.tolic,damjan.miklic,zdenko.kovacic}@fer.hr
[2] Faculty of Education and Rehabilitation Sciences, University of Zagreb, Zagreb, Croatia
{mcepanec,s.simlesa}@gmail.com

Abstract. The existing procedures for Autism Spectrum Disorder (ASD) diagnosis are often time consuming and tiresome both for highly-trained human evaluators and children. In addition, prospective human evaluators need to undergo a rigorous and lengthy training process that may not be accessible or affordable to all interested individuals. Hence, this paper proposes a framework for robot-assisted ASD evaluation based on Partially Observable Markov Decision Process (POMDP) modelling. POMDP is broadly used for modelling optimal sequential decision making tasks under uncertainty. Spurred by the widely accepted Autism Diagnostic Observation Schedule (ADOS), we start off with emulating ADOS. In other words, our POMDP model explicitly takes into account the ADOS stratification into several modules, ongoing task informativeness and robotic sensor deficiencies. Relying only on imperfect sensor observations, the robot provides an assessment of the child's ASD-relevant functioning level (which is partially observable) within a particular task. Finally, we demonstrate that the proposed POMDP framework provides fine-grained outcome quantification, which could also increase the appeal of robot-assisted diagnostic protocols in the future.

Keywords: Robotics · POMDP · Autism spectrum disorder · Diagnostics

1 Introduction

Autism Spectrum Disorder (ASD) is a developmental disorder characterised by impairment in social interaction, verbal and nonverbal communication and by repetitive behaviours and interests. It has become a commonly diagnosed neurodevelopmental disorder, with increasing prevalence rates, affecting between 1.1% and 1.5% of children population [1]. With an estimated yearly prevalence

H. Liu et al. (Eds.): ICIRA 2015, Part II, LNAI 9245, pp. 82–94, 2015.
DOI: 10.1007/978-3-319-22876-1_8

increase of 3%, it is expected to become one of the most common disorders in the near future. To increase the inclusion rate and reduce the cost of lifelong care for people with ASD, experts focus on early diagnostics and intervention.

Since there are no medical markers of autism that could be used in a diagnostic process, the diagnosis relies on behavioural observations made by experienced clinicians. There are several mutually non-exclusive approaches to diagnostics: a) using criteria from the Diagnostic and Statistical Manual of Mental Disorders (DSM) [2]; b) testing children using the Autism Diagnostic Observation Schedule (ADOS) [3]; and, c) interviews with the caregivers using the Autism Diagnostic Interview-Revised (ADI-R) [4].

Studies show that the agreement between clinicians on different DSM criteria for autism varies from 0.58 to 0.79 [5]. For ADOS, the inter-rater reliability for some ratings is as low as 0.38 in modules involving preschool children [3]. The main reason for these discrepancies is that the diagnostic procedure is highly complex due to simultaneous observation, coding and interpretation of many behaviours as well as administration of various specific tasks. Additionally, the process of learning to observe and code the behavior in order to achieve 80% inter-rater reliability on ADOS might last for several years. Moreover, the process of diagnosing and assessing the degree of disorder can take several years even for experienced clinicians, delaying the all-important early intervention. Thus, there exists a need for a more objective approach that would help clinicians in gathering multimodal information and coding the social behavior. Modern robotics technologies seem capable of providing adequate tools to address this need.

Since children with ASD are often attracted by technological devices, there are many robotic applications focusing on teaching and intervention [6,7]. However, diagnostic applications are scarce even though *there exists a need for quantitative, objective measurements of social functioning for diagnosis, for evaluating intervention methods, and for tracking the progress of individuals over time* [8]. This is mainly due to the complexity of the diagnostic process itself, but is also dependent upon efficient processing and reasoning algorithms that are to be implemented on the robot. The work in [9] proposes several quantitative metrics for social response during the diagnostic process, with data collected through passive sensors installed in the examination room. In addition, there is an obvious trend in computer science to improve the diagnostic process for ASD (although robots are not deployed). Such an effort is reported in [10], where language and speech are processed in order to quantify some of the behavioural patterns. Through observation of natural conversation between the examiner and child, it is shown that the behavior of examiners changes depending on the perception of the child's behavior, indicating the need for consistent and repetitive stimuli (also called *social presses*), which can be achieved through the use of a robotic examiner.

The present work is the first step towards devising an adaptive protocol aimed at expediting the diagnostics procedure while retaining (and possibly increasing) the reliability of the standard ADOS protocol. This implies the robot [11] being able to act and react under uncertainty and infer information about the child from observing the behaviour, similarly to human evaluators. Due to their generality and similarity to the human decision-making process under uncertainty, we adopt Partially Observable Markov Decision Processes (POMDPs) to model the interaction between the robot and children. POMDPs are already used in a variety of applications for human-robot interaction and acting under uncertainty. In [12], the authors explore the model of interaction between a disabled person and a robotic wheelchair, while the authors of [13] model humans as observation providers to a robotic system navigating in unknown environments. Similarly to our concept of a robot eliciting the interaction in order to obtain information about the child, [14] utilizes POMDPs to improve robot plans for information gathering in unknown environments. Moving away from navigation applications, POMDPs are also used to model human-robot dialogue [15] and human-in-the-loop systems [16].

Herein, we primarily focus on the POMDP model of diagnostic tasks and use the belief state of the model as an automatic evaluator of the child's behaviour. In other words, we do not yet fully exploit the POMDP decision-making feature, which is the subject of our future work. Therefore, the contributions of this paper are twofold: a) the ASD diagnostic task modelling via Partially Observable Markov Decision Process (POMDP); and, b) the automatic coding of the child's behaviour at the end of each task through the POMDP belief state.

The paper is organized as follows. Section 2 presents two tasks of the robot-assisted ASD diagnostic protocol along with general information about POMDPs. In Section 3, we model the two protocol tasks employing POMDP, which automatically codes the child's behaviour through the POMDP belief state. A comparison with the human evaluator is in Section 4. Concluding remarks and future directions are given in Section 5.

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 behaviour. The principal question is whether the robot can help clinicians reach a diagnosis reliably in shorter time. Our robot-assisted ASD diagnostic protocol consists of four tasks developed upon the ADOS protocol and is aimed for children aged 2-6. In this paper we focus on two tasks: *response to a name call* and *joint attention* from our previous work [17]. During the clinical sessions it was observed that the proposed tasks, when performed by the robot, take significant amount of time to complete. Therefore we shorten the tasks in this work.

The *response to a name call* 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 progress of the task is shown in Fig. 1. At the beginning of the task, the child is distracted by playing with some object. Meanwhile, the robot calls (action *call*) the child by name and repeats this action at most four times if the child does not respond. Once the child responds, the task terminates (action *end*). If the child does not respond to being called by name, the robot reinforces the call by referencing the child's favourite toy (action *rcall*), food or activity. After the reinforced call, the task is completed, as presented in Fig. 1 in the form of a decision tree.

The *joint attention* task of the protocol focuses on a child's ability to transfer attention from one object of interest to another. The positive response is again verified via eye contact. In this task, one robot calls the child and turns its head towards the other robot (action *turn*) trying to make the child transfer its attention to the other robot. This turning procedure is repeated at most four times if the child does not respond, followed by a reinforced call with pointing towards (action *point*) the other robot. If the child does not respond to the pointing gesture, the other robot, which was passive up to this point, tries to attract the child's attention by waving its arms, flashing LEDs or producing various sounds (action *attract*). After this action, the task ends (action *end*) (see Fig. 1).

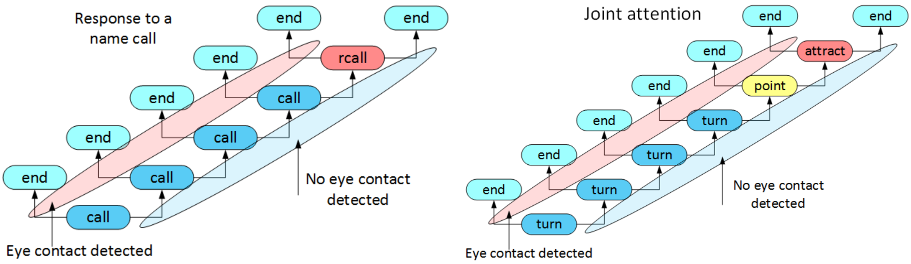

Fig. 1. Decision trees for *response to name* and *joint attention*

One can notice a salient pattern in these tasks. Namely, the sequence of performing an action, observing the child's response and selecting the next action based on the results of the previous one while building the opinion on the state of the child. Notice that the child state, as far as ASD is concerned, cannot be observed directly. This cyclic pattern fits the framework of POMDPs, which we use to model the tasks of the robot-assisted ASD diagnostic protocol.

2.1 Partially Observable Markov Decision Processes

A POMDP is defined as a tuple (S, A, O, T, Ω, R), where S is a set of hidden states of the process, A is a set of actions that can be performed, O is a set of observations generated by the process, 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 process is hidden, we maintain 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 system chooses the 'best' action with respect to the expected reward. This mapping from belief state to the set of actions is called policy and is denoted π.

After an action a is taken, the process transitions from a state s to a state s' and generates an observation o' (the process is also allowed to stay in the same state, that is, $s' = s$). At this time, the belief state is updated using the Bayes' rule as follows:

$$b'(s) = p(s'|a, b, o') = \frac{p(o'|s', a) \sum_{s \in S} [p(s'|s, a) \cdot b(s)]}{p(o'|a, b)}. \quad (1)$$

Solving a POMDP yields an optimal policy π such that the expected cumulative reward is maximized. Finding an exact solution to a POMDP is usually deemed intractable. Hence, there exist approximate solvers such as *pomdp-solve* [18] or *SymbolicPerseus* [19].

3 Modelling Diagnostic Protocol Tasks via POMDPs

In the proposed robot-assisted diagnosis, the robot evaluator builds its opinion about the state of the child and presents it to expert evaluators, along with the log of all data. Therefore, the POMDP model incorporates all states of a child that are of interest in the assessment process, which belong to a set S_c, while the robot builds its belief b over the states $s_c \in S_c$. This belief b reflects the similarity between the child's observed behaviour and the child's expected behaviour when in a state s_c. We focus on three states related to children having low functioning ASD, high functioning ASD and children not suffering from ASD, which we denote $S_c = \{s_{LA}, s_{HA}, s_{NA}\}$, respectively.

The main social cue, on which we focus herein, is eye contact. Therefore, we keep track of the following two observations: $o_c \in O_c := \{o_{EC}, o_{NEC}\}$, meaning eye contact detected and no eye contact detected, respectively. Child states S_c and tracked observations O_c comprise the core of the task POMDP model and are found in both tasks we focus on. Let us now model the specifics of each task.

3.1 Response to a Name Call

Response to a name call consists of four calls by name followed by one call that uses a special phrase referring to something very dear or interesting to the child. According to Fig. 1, the task ends when the child establishes eye contact after a call. 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 most important parts of POMDP models are state transition and observation probabilities. These probabilities are often extracted from a large amount of data via learning algorithms (Hidden Markov Models being a prominent example). However, since there are no large amounts of data related to the ASD diagnostics available, we approach modelling more pragmatically. Since the robot actions cannot change the aforementioned states, we define the following state transition probabilities for any given action:

$$\begin{aligned} p(s'_{LA} \mid s_{LA}, a \in A) = 1, \\ p(s'_{HA} \mid s_{HA}, a \in A) = 1, \\ p(s'_{NA} \mid s_{NA}, a \in A) = 1, \end{aligned} \quad (2)$$

while all other conditional probabilities of transitions between the states in S_c are set to 0.

To the best of our knowledge, there are no studies from which observation probabilities could be inferred. As far as the ADOS Module 2 is concerned, the following model parameters for action *call* have been gathered from expert clinicians from the Faculty of Education and Rehabilitation Sciences, University of Zagreb:

$$p(o_c = o_{EC}) = \begin{cases} 0.8, & \textit{children without ASD}, \\ 0.5, & \textit{children with high functioning ASD}, \\ 0.2, & \textit{children with low functioning ASD}. \end{cases} \quad (3)$$

The reinforced call results in a higher probability of a child to establish eye contact, which we model with a slight increase in the observation probability with respect to (3), that is,

$$p(o_c = o_{EC}, a = rcall) = \begin{cases} 0.9, & \textit{children without ASD}, \\ 0.6, & \textit{children with high functioning ASD}, \\ 0.3, & \textit{children with low functioning ASD}. \end{cases} \quad (4)$$

The observation probability distribution is set to uniform for the action *end*.

POMDP optimal policies can be sought over an infinite number of actions or over a finite horizon. However, in the case of our diagnostic tasks, the number of actions taken is limited. The goal is to obtain as much information as possible without exhausting the child or making it bored and uninterested. To track the progress of the task and the amount of information gathered, we introduce a new set of states, $s_I \in S_I := \{s_{LI}, s_{HI}\}$, denoting low information state and high information state, respectively. Thus, the states of the POMDP are $S = S_c \times S_I$.

We model the informativeness of an action $a \in A$, denoted I_a, as the probability of a transition from state s_{LI} to state s_{HI}:

$$I_a = p(s_{HI} \mid s_{LI}, a \in A). \quad (5)$$

The robot cannot lose information when performing an action; hence:

$$p(s_{HI} \mid s_{HI}, a \in A) = 1, \qquad p(s_{LI} \mid s_{HI}, a \in A) = 0. \quad (6)$$

Since transitions between information states in S_I do not depend on the child states in S_c and observations in O_c, we deduce the following belief update formula for information states by extracting the constant parts from (1):

$$b'(s_I) = k \sum_{S_I} \left[p(s'_I | s_I, a \in A) \cdot b(s_I) \right]. \quad (7)$$

The belief over states S_I must be normalized, i.e., $\sum b(S_I) = 1$, which means that k can be omitted from expression (7), yielding:

$$b'(s_I) = \sum_{S_I} \left[p(s'_I | s_I, a \in A) \cdot b(s_I) \right]. \tag{8}$$

The robot is tracking the amount of information gathered by updating the belief $b(s_{HI})$. Plugging (5) and (6) into (8), the following holds:

$$b'(s_{HI}) = b(s_{HI}) + I_a \cdot [1 - b(s_{HI})]. \tag{9}$$

We now constrain the number of actions in the task by choosing I_a such that, within the desired horizon of n steps, the belief $b(s_{HI})$ reaches 0.9. The task is then terminated by performing the action $a = end$, which is ensured by highly rewarding it when the amount of information collected is high, and by penalizing it when the amount of information is low. It is also important to note that information update (9) results in subsequent actions being less and less informative. Besides avoiding child fatigue, this property is the reason why diagnostic tasks involve a finite (and relatively small) number of repetitions.

Since the *response to a name call* task requires 5 actions prior to completion (not counting the ending action), we model informativeness of actions as $I_a = 0.4$, for which $b(s_{HI})$, updated according to (9), attains a value greater than 0.9 after 5 actions. Note that all actions are considered equally informative.

The main obstacle when using this approach pertains to how the solver handles the ending action. Essentially, the value iteration of POMDP solvers does not account for the possibility that no more actions will be performed after the action *end* is performed. Therefore, we extend the above model by adopting the approach similar to the tiger problem [20], with the *end* action resetting the belief state b to the uniform distribution over the state set S_c. This is achieved by defining both the transition and observation probability distribution as the uniform distribution over S_c for the action *end*. With $p(s'|s, end) = const$ and $p(o'|s', end) = const$ for all $s' \in S_c$, the following holds:

$$\begin{aligned} b'(s) &= \frac{p(o'|s', end) \sum_{s \in S_c} p(s'|s, end) b(s)}{p(o'|end, b)} = \\ &= \frac{p(o'|s', end) p(s'|s, end) \sum_{s \in S_c} b(s)}{p(o'|end, b)}. \end{aligned} \tag{10}$$

Since $\sum_{s \in S_c} b(s) = 1$ and taking into account that $p(o'|end, b)$ is the normalizing factor used to ensure $\sum_{s \in S_c} b'(s) = 1$, we infer from (10) that $b'(s)$ is constant for all $s \in S_c$. Equivalently, the belief state b' is reset to the uniform distribution.

In addition to the reset of belief over states S_c, resetting the belief over information states S_I also has to be performed. In this case, the next iteration of the task starts from $b(s_{HI}) = 0$, which is achieved by setting $p(s_{HI}|s_I, end) = 0$. In other words, the process cannot be in the state s_{HI} after the action *end*, regardless of the observation.

Since the belief state is reset after the action *end*, an optimal policy between two ending actions remains the same and is interpreted as a robot performing the same task over and over again. Therefore, the limited amount of actions is successfully modelled, while letting the solver iterate over multiple instances of a task to generate the optimal policy.

The reward values for states and are manually tuned to obtain the policy in Fig. 1, following the rules:

- if $b(s_{HI})$ is low, reward *call*;
- if $b(s_{LI}) + b(s_{HI})$ is high, reward *rcall*; and
- if $b(s_{HI})$ or $b(s_{NI})$ is high, reward *end*.

It is important to note that multiple reward models can generate the desired policy.

3.2 Joint Attention

We adopt a similar approach to the model of the *joint attention* task. Sets of states $S = S_c \times S_I$ and observations O_c remain the same as for the *response to a name call* task. Additionally, we keep the property that actions do not change the state s_c, thus defining transition probabilities the same as in (2). Since the *joint attention* task consists of four calls accompanied by turning the head, one call accompanied by pointing and one attempt of attracting the attention of the child, the set of actions for this task is $A = \{turn, point, attract, end\}$. The following observation probabilities for given actions have been provided by experienced clinicians:

$$p(o_{EC}, turn) = \begin{cases} 0.7, & \textit{children without ASD}, \\ 0.4, & \textit{children with high functioning ASD}, \\ 0.2, & \textit{children with low functioning ASD}, \end{cases} \tag{11}$$

$$p(o_{EC}, point) = \begin{cases} 0.8, & \textit{children without ASD}, \\ 0.5, & \textit{children with high functioning ASD}, \\ 0.3, & \textit{children with low functioning ASD}, \end{cases} \tag{12}$$

$$p(o_{EC}, attract) = \begin{cases} 0.9, & \textit{children without ASD}, \\ 0.6, & \textit{children with high functioning ASD}, \\ 0.4, & \textit{children with low functioning ASD}, \end{cases} \tag{13}$$

with the observation probability distribution for the action *end* set to uniform. Since this task requires 6 actions to be completed, we model the informativeness of actions as $I_a = 0.35$, while we tune the rewards to achieve the policy described in Fig. 1.

Optimal action graphs for both tasks are presented in Fig. 2, showcasing that it is possible to emulate the existing human reasoning by using the proposed POMDP models and carefully designing rewards. In Fig. 2 we present the sequence of actions that are to be performed for the *response to a name call* task. The belief $b(ASD)$ is the combined belief over states of low and high functioning

ASD. The robot starts with the uniform belief b over S_c; thus, $b(ASD) = 0.667$. The action to be taken at this point is *call*, after which the belief $b(ASD)$ branches depending on the observation. Based on the updated belief, the robot chooses the optimal action in the second iteration. In the showcased scenario, if the child responded the robot's action is *end* (left branch), while it continues to *call* the child in case of no positive response (right branch). A similar reasoning is applied to the rest of the task and the *joint attention* task.

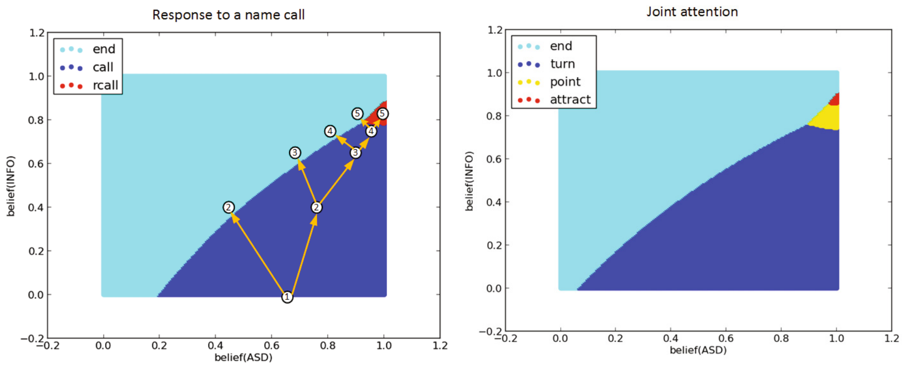

Fig. 2. Optimal actions with respect to the belief state

3.3 Modelling Sensor Accuracy

Due to the complexity of the behaviours tracked by the robot, perfect accuracy of the perception algorithms cannot be attained. The performance of perception algorithms is accounted for by modelling the sensors' accuracy in the form of uncertainty, which allows for easy integration into the POMDP model. To that end, we define the accuracy of the sensor as a conditional probability p_s:

$$p_s = p(\hat{o_c}|o_c), \tag{14}$$

where $\hat{o_c}$ is the observation perceived by the sensor, while o_c is the observation generated by the process (i.e., the child). The conditional observation probabilities of POMDP model are transformed by taking into account the probability of detection of the real observation and probability of detection of false positives:

$$\begin{aligned} p'(o_c|s_c, a) &= p_s \cdot p(o_c|s_c, a) + [1 - p_s] \cdot \sum_{o \neq o_c} p(o|s_c, a) = \\ &= p_s \cdot p(o_c|s_c, a) + [1 - p_s] \cdot [1 - p(o_c|s_c, a)] \,. \end{aligned} \tag{15}$$

4 Performance of the POMDP Modelled Diagnostic Protocol

The tasks in the ADOS protocol are coded with numbers $(0, 1, 2, 3)$ independently, then the results of individual tasks are summarized and a threshold is applied to decide whether the child has ASD or not. The coding itself takes place both simultaneously with the diagnostic session, but also takes a large amount of time after the session, sometimes even including the review of the video recording of the session.

We propose to encode the tasks of robot-assisted diagnostic protocol with the belief state b obtained from the POMDP model of each task, maintaining a scoring scheme similar to that of ADOS, but with finer granulation, possibly enabling the simultaneous assessment of the degree of disorder. The codes for the *response to a name call* task are summarized in Table 1. Each code corresponds to the belief $b(s_c)$ when the *end* action is executed. As already shown in Fig. 1, there are 6 possible outcomes for this task.

Table 1. Codes for *response to a name call* with different sensor accuracies p_s

State	p_s	Code $= b(s_c)$					
$s_c = s_{LA}$	100%	0.13	0.28	0.45	0.60	0.76	0.92
	70%	0.25	0.33	0.41	0.48	0.57	0.70
$s_c = s_{HA}$	100%	0.33	0.44	0.44	0.36	0.23	0.08
	70%	0.33	0.35	0.35	0.33	0.31	0.24
$s_c = s_{NA}$	100%	0.54	0.28	0.11	0.04	0.01	0
	70%	0.42	0.32	0.24	0.19	0.12	0.06

Comparing the codes for different states, we infer that $b(s_{NA})$ is the highest when the child responds immediately, as expected. If the child responds within the first several iterations of social presses, the similarity to the behaviour of children with high functioning ASD is the highest. With the task extending into more than 3 robot actions, $b(s_{LA})$ becomes the dominant component of belief state b, indicating that the behaviour is very similar to that of a child with low functioning ASD. We also explore the values of b when we introduce a non-perfect vision sensor. The values of belief states are similar to those obtained by the perfect sensor, while the differences between the belief over states for an individual outcome are smaller, which is expected since with a lower accuracy the observation provides less information about the process state. Similar conclusions are drawn regarding the *joint attention* task (see Table 2).

In addition to exhibiting similar trends, the codes for both tasks are comparable with respect to the task outcome, which is in line with the current scoring scheme where each task has the same weight in the final result. Using codes from Tables 1 and 2, taking into account that we estimate our eye contact detection

Table 2. Codes for *joint attention* with different sensor accuracies p_s

State	p_s	Code = $b(s_c)$						
$s_c = lowASD$	100%	0.15	0.26	0.38	0.49	0.63	0.74	0.87
	70%	0.26	0.32	0.38	0.44	0.5	0.56	0.66
$s_c = highASD$	100%	0.30	0.39	0.43	0.42	0.33	0.25	0.13
	70%	0.33	0.34	0.35	0.35	0.34	0.32	0.28
$s_c = noASD$	100%	0.54	0.35	0.19	0.09	0.04	0.01	0
	70%	0.41	0.34	0.27	0.21	0.16	0.11	0.06

algorithm to be 70% accurate, we encode the results of our previous sessions with children [17] using $b(s_{LA})$ as code. The sessions were performed with three children with ASD (labelled ASD01-ASD03) and one typically developing child (labelled TYP01). The codes obtained by the proposed POMDP models are compared with the codes obtained via the traditional ADOS scoring scheme and presented in Table 3. Note that the codes obtained via POMDP, which belong to the interval $[0, 1]$, are scaled to ADOS interval, which is $[0, 3]$.

Table 3. Comparison of codes for behaviours observed in sessions with children

Task	Child	Response (iteration)	Code: POMDP	Code: scaled POMDP	Code: ADOS
Response to a name call	ASD01	4	0.48	**1.44**	**1**
	ASD02	3	0.41	**1.23**	**1**
	ASD03	No	0.70	**2.10**	**3**
	TYP01	No	0.70	**2.10**	**3**
Joint attention	ASD01	No	0.66	**1.98**	**3**
	ASD02	No	0.66	**1.98**	**3**
	ASD03	6	0.56	**1.68**	**2**
	TYP01	∅	∅	∅	∅

Comparing the scaled POMDP codes with ADOS codes (the last two columns in Table 3), we conclude that the POMDP scoring scheme possesses finer resolution. Namely, ADOS assigns the code 1 both to ASD01 and ASD02 while our POMDP code distinguishes these two children. We expect this finer resolution property to prove beneficial for the adaptive protocol we aim to develop.

5 Conclusion

Motivated by the limitations of existing autism diagnostic protocols, we propose a framework for robot-assisted ASD evaluation based on POMDP modelling. The proposed POMDP design allows the user to explicitly model the targeted participant population, ongoing task informativeness and robotic sensor deficiencies. Our framework utilizes the POMDP belief state to quantify the functioning ASD level. We also show that, by thoughtfully designing the reward function of the POMDP, one can emulate the reasoning of ADOS protocol, while providing higher resolution of the outcome by using the belief states generated via our POMDP models.

These favourable numerical results are reassuring and prompt an extensive clinical trial involving several dozens of participants. Next to providing a hands-on verification of the work presented herein, this clinical trial will also allow us to fine tune the POMDP parameters. Afterwards, our goal is to increase the resolution of the assessment by tracking more sets of observations, such as vocalizations and gestures which are already tracked within ADOS. Ultimately, we plan to further exploit the decision-making feature of POMDPs and design an adaptive diagnostic protocol that could conceivably accelerate the diagnostic procedure and make it more reliable.

Acknowledgments. This work has been fully supported by Croatian Science Foundation under the project Autism Diagnostic Observation with Robot Evaluator (no. 93743-2014).

References

1. Baio, J.: Prevalence of autism spectrum disorder among children aged 8 years - autism and developmental disabilities monitoring network, 11 sites, United States, 2010. MMWR **63**(2) (2014)
2. Diagnostic and Statistical Manual of Mental Disorders, 4th edn. American Psychiatric Association (2002)
3. Lord, C., Rutter, M., Dilavore, P.C., Risi, S.: Autism Diagnostic Observation Schedule. Western Psychological Services (2002)
4. Rutter, M., LeCouteur, A., Lord, C.: The Autism Diagnostic interview, Revised (ADI-R). Western Psychological Services (2003)
5. Klin, A., Lang, J., Chicchetti, V., Volkmar, F.R.: Interrater reliability of clinical diagnosis and DSM-IV criteria for autistic disorder: results of the DSM-IV autism field trial. Journal of Autism and Developmental Disorders **30**(2), 163–167 (2000)
6. Iacono, I., Lehmann, H., Marti, P., Robins, B., Dautenhahn, K.: Robots as social mediators for children with autism - a preliminary analysis comparing two different robotic platforms. In: 2011 IEEE International Conference on Development and Learning (ICDL), vol. 2, pp. 1–6 (2011)
7. Liu, C., Conn, K., Sarkar, N., Stone, W.: Online affect detection and robot behavior adaptation for intervention of children with autism. IEEE Transactions on Robotics **24**(4), 883–896 (2008)
8. Scassellati, B., Henny, A., Mataric, M.: Robots for use in autism research. Annual Review of Biomedical Engineering **14**(1), 275–294 (2012)
9. Scassellati, B.: Quantitative metrics of social response for autism diagnosis. In: IEEE International Workshop on Robot and Human Interactive Communication, ROMAN 2005, pp 585–590 (2005)

10. Narayanan, S., Georgiou, P.G.: Behavioral Signal Processing: Deriving human behavioral informatics from speech and language. Proceedings of the IEEE **101**(5), 1203–1233 (2013)
11. Aldebaran Robotics. Nao Software Documentation, v2.1 edn. (2015). https://community.aldebaran-robotics.com/doc
12. Taha, T., Miro, J.V., Dissanayake, G.: A pomdp framework for modelling human interaction with assistive robots. In: 2011 IEEE International Conference on Robotics and Automation (ICRA), pp. 544–549, May 2011
13. Rosenthal, S., Veloso, M.: Modeling humans as observation providers using pomdps. In: RO-MAN 2011, pp. 53–58. IEEE, July 2011
14. Wang, M., Canu, S., Dearden, R.: Improving robot plans for information gathering tasks through execution monitoring. In: 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 5285–5291 (2013)
15. Caccavale, R., Leone, E., Lucignano, L., Rossi, S., Staffa, M., Finzi, A.: Attentional regulations in a situated human-robot dialogue. In: The 23rd IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN 2014, pp. 844–849, August 2014
16. Lam, C.-P., Sastry, S.S.: A pomdp framework for human-in-the-loop system. In: 2014 IEEE 53rd Annual Conference on Decision and Control (CDC), pp. 6031–6036, December 2014
17. Petric, F., Hrvatinic, K., Babic, A., Malovan, L., Miklic, D., Kovacic, Z., Cepanec, M., Stosic, J., Simlesa, S.: 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, October 2014
18. Cassandra, A.R.: pomdp-solve, v5.3 edn. (2013). http://pomdp.org/
19. Poupart, P.: Exploiting Structure to Efficiently Solve Large Scale Partially Observable Markov Decision Processes. PhD thesis, University of Toronto (2005)
20. Cassandra, A.R., Kaelbling, L.P., Littman, M.L.: Acting optimally in partially observable stochastic domains. In: Proceedings of the 12th National Conference on Artificial Intelligence, pp. 1023–1028 (1994)