A POMDP Design Framework for Decision Making in Assistive Robots

Conference Paper · May 2017

DOI: 10.1007/978-3-319-58071-5_35

4 authors, including:

Ioannis Kostavelis
The Centre for Research and Technology, Hellas

Dimitrios Giakoumis
The Centre for Research and Technology, Hellas

Dimitrios Tzovaras
The Centre for Research and Technology, Hellas

A POMDP Design Framework for Decision Making in Assistive Robots

Ioannis Kostavelis, Dimitrios Giakoumis, Sotiris Malassiotis, and Dimitrios Tzovaras

Centre for Research & Technology Hellas, Information Technologies Institute, 6th Km Charilaou-Thermi Road, 57001, Thermi-Thessaloniki, Greece
{gkostave,dgiakoum,malasiot,tzovaras}@iti.gr

Abstract. This paper proposes a theoretical framework that determines the high-level cognitive functions for multipurpose assistive service robots, required to autonomously complete their tasks. It encompasses a probabilistic POMDP based decision-making strategy that provides constant situation awareness about the human and the environment by associating the robot awareness about the user with specific clusters of robotic actions. To achieve this, a method for designing POMDP models is presented herein ample to define decision making policies suitable to resolve assistive tasks through a series of robotic actions. The proposed POMDP design methodology compensates the partial and noisy sensor input acquired from the robot sensors by foreseen mitigation strategies on the robot's decisions when a software component fails. The theoretical work presented herein is assessed over well defined robotic tasks and proved capable to operate in realistic assistive robotic scenarios.

Keywords: decision making, assistive robots, POMDP, partial observability, robot situation awareness

1 Introduction

The autonomy of the contemporary assistive robots relies on their capacity to decide on their own actions based on their cognitive functionalities and realize these actions using their robotic planning mechanisms [1]. However, robot autonomy is not an end in itself in the field of Human Robot Interaction (HRI) but rather a way to support bidirectional interaction using robot actuator movements, communication and representational schemas [2]. Based on this statement, the modelling procedure of the robot cognitive functions, which will eventually determine the decision making mechanism, should take into consideration that the realization of an inferred robotic action will alter the status of the environment and the interaction schema with the user, making thus the robot an active participant in the human-robot cohabitation rather than a passive observer. Following this notion, Markov Decision Processes (MPDs) constitute an efficient solution for decision making and have beeb proved adequate to solve simplified problems with diminished uncertainty [3]. However, it is typical for humans in real life scenarios to make decisions under uncertainty since not all the facts are measurable and not all the required observations are constantly available. When it comes to robotic application development, complete situation awareness is not feasible since the environment should be modeled with limited robot sensors and the acquired sensor observations are noisy; in such applications, the robot belief uncertainty about the current state of the human, the environment and the robot itself is broadened. Partially Observable Markov Decision Processes (POMDPs) are able to model the uncertainties stemming from realistic situations better than the MPDs, while their main difference is that the world's state in the POMDP is not known to the robot; instead, a probabilistic observation corresponding to the state is received from the environment after performing each action [4].

The paper at hand aims to model the high-level cognitive functions of an assistive social robot by formulating a decision making mechanism based on POMDP models. Specifically, an explicit methodology for designing POMDP models is presented herein, comprising a generalized theory in the domain of robot intervention in assistive tasks. This method capitalizes on robot's various perception modalities, yet partially available, to infer an action plan considering clustered type of robotic actions. Contrary to the existing POMDP applications where the designed POMDP models are constrained to resolve specific tasks by selecting optimal robotic actions, the proposed solution tackles the problem from an alternative view point, where the selection of the robotic actions is tightly related to the robot's awareness about the human condition, which is reflected into robot's alert levels. Subsequently, the task to be resolved is the propagation of the robot's states to a lower level of alert.

2 Related Work

2.1 POMDPs as prompting systems

A profusion of laborious research has been conducted in the filed of health technology by introducing systems for elderly, possibly with cognitive or physical disabilities, who want to continue living independently in their own homes [5]. Under this scope POMDPs have been either utilized as prompting systems assisting people in their daily life by verbally motivating them to successfully complete specific activities [6], or integrated in robotic agents modelling their action planning mechanism to fulfill their assistive task. The design of prompting systems that react on time while exploit sensing and modelling mechanisms is a laborious work. One exemplar application is the COACH system [7] which uses computer vision to monitor the progress of a person with dementia washing her/his hands and prompts only when necessary by employing a POMDP acting as a temporal probabilistic model based on the sensed observations. However, the COACH system is tailored to specific tasks and requires great amount of expert knowledge for re-designing the POMDP model in order to be useful for different tasks with generalization capacities. The authors in [8] tried to introduce a more generalized framework for building prompting systems using POMDP models by incorporating psychologically justified expert background knowledge. Specifically this method incorporated Interaction Units [9] which is a psychologically justified description of the task and the particular environment where this task is to be carried out that can be generated from empirical data. This is then combined with a specification of the available sensors to build a working prompting system based on POMDPs. However, the automation of this procedure for the production of context aware POMDP models is also limited and still requires the expertise of a psychologist. As an attempt to diverge from the phycological modeling the authors in [10] proposed a probabilistic relational model encoded in a relational database allowing non experts in POMDP design to fill in the necessary details of a task using a simple and intuitive procedure. Although this method proved capable to automatically produce POMDP models for assistive applications, yet the probabilistic database framework was limited to the scale of the problem that could be modeled, also deteriorating the prompting capacities to specific tasks that require well restricted operational environment.

2.2 POMDPs in robotic applications

In the service robots domain, the POMDPs have been widely utilized to increase the robot autonomy in human populated environments. This has been achieved in multiple applications levels, concerning examples such as navigation and manipulation. The first attempt for using POMDPs in robotic tasks is the work discussed in [11], where action space was simplified in basic and discretized moves of the robot while the received observations were abstract representations of the environment. The authors in [12] solved a more complicated aspect of the navigation problem using POMDPs, where a point robot was considered by diminishing the state and action space accordingly. The authors proved that a simplified action space during the POMDP design can closely resemble the efficacy of a model with greater resolution in the action space, while also the policy computation time was discussed. In a more sophisticated work, the authors in [13] utilized a hierarchical POMDP framework for robot navigation in the context of which the localization, the local planning and obstacle avoidance was tackled. In this work the occupancy grid map comprised the state space of the robot, while the actions space consisted of a hierarchical discrimination of robot rotational and translation capabilities. In [14], a POMDP model has been developed to determine an objective function that considers both probability of collision and uncertainty at the goal position, providing an alternative path planning decision policy. Moreover, a proof of the usage of POMDPs in a great variety of robotic applications is the work described in [15]. In this application, the authors utilized a POMDP model as a prompting system in a probabilistic planning of a bimanual robot that was targeted to unfold clothes. At this stage, it can be inferred that although the aforementioned methodologies were addressing their targeted functionalities competently, they were limited to their specific task. This can be partially explained by the fact that the computational complexity of solving a POMDP problem instance grows exponentially with the size of the state and action space and thus it is difficult to concurrently model precisely the state and the action space for the human, the robot and the environment. Therefore, specific works tackle the issue from a different view point where the POMDPs are utilized as robot control mechanisms that orchestrate the robots behaviour in a variety of application tasks. In the work introduced in [4], the authors developed hierarchical POMDP models focusing in the abstraction of the action space linking the hierarchy levels with actions tightly related with folded subordinate POMDP models. In this way the authors achieved to design easily manageable smaller POMDP models dedicated to specific robotic actions instead of modeling a global model which is hardly to be solved by the existing POMDP solvers [16]. This way, multiple tasks can be modeled using the POMDP theory. In a more contemporary work [17], the authors introduced a decision making and control supervision system suitable to operate on multi-modal service robots. This bridges the gap of abstraction between designed POMDP models and the physical world concerning actions, while multi-modal perception is processed to extract measurements uncertainty. Complementary to the aforementioned works, a method that determines the human robot interaction with assistive robots through POMDPS is the one described in [18]. The authors designed a POMDP model where the human satisfaction from the collaboration with the robot is the key factor to model the interaction, while the status of the user in terms of awareness and stress determines the human's participation in the execution of the task. However, during the design of the POMDP model, the target task should be explicitly analysed in the expected states of both the human and the robot, comprising an efficient yet hard to be modeled solution.

Considering the existing solutions, the added value of the proposed work is the determination of a theoretical framework that describes an explicit POMDP design methodology. The formulated POMDP models are human-centric and drive the cognitive functions of multipurpose social assistive robots. This statement is based on the principal that the robot should be always on alert and aware about the human cohabitant, while the decided actions should reflect the level of robot's alert. Through this procedure, the robotic actions are tightly related to the amount of assistance that is required to offer to the cohabitant. Specifically, when the robot's level of alert about the human is increased, the planned robot actions should be more intensive and interventional in order the robot to become complacent about the human, while in intermediate levels of robot alert more discreet actions should be planned.

3 Proposed Method for Decision Making Design

3.1 Robotic-wise POMDP formulation

For the proposed problem formulation, it is essential to interpret the generic POMDP design theory [19] in a robotic-wise manner, considering the explicit assistive robot scenarios where the problem domain comprises the environment, the human and the robot. Towards, this direction, the discrete POMDP is designed as a tuple $P = \{S, A, \Omega, R\ O, T, b_0\}$ where:

- $S = \{s_1, s_2, ..., s_n\}$ denotes the **States** space that determines the condition of the environment, the human and the robot at each time t.
- $A = \{a_1, a_2, ..., a_n\}$ denotes the **Actions** space that encloses all the actions that the robot is able to perform so as to interact with the human and the environment.

- $\Omega = \{\omega_1, \omega_2, ..., \omega_n\}$ denotes the **Observations** space that comprises the robot perception input from the human and the environment, yet under the assumption that an observation ω partially describes the state of the previous entities.
- $R = (A, S)$ comprises a **Reward** function that determines the restrictions imposed by penalizing or endorsing specific robotic actions (A) during the interaction with the human and the environment (S).

The challenging part during the design of POMDPs is the determination of the probability distributions of the initial state (b_0), the states transitions (T) and the observations (O), something that makes the POMDP designers to frequently rely on phycologists to empirically quantify these values under assumptions that are difficult to be assessed.

- The probability distribution of the *initial state* comprises the likelihood about the environment, the human and the robot to be in specific state s at the time $t = 0$ such as:
$$b_0(s) = P(s_0 = s) \tag{1}$$
- The probability distribution of the *state transition* comprises the probability of propagating to state s' given that the domain is in state s and the robot selects an action a and the its respective expression is provided as:
$$T(s, a, s') = P(s_t = s' | s_{t-1} = s, a_{t-1} = a) \tag{2}$$
- The probability distribution of the *observations* comprises the uncertainty for the perception of an observation ω considering that the environment and the human are in state s and the robot has performed the action a, also expressed as:
$$O(s, a, \omega) = P(\omega_t = \omega | s_{t-1} = s, a_{t-1} = a) \tag{3}$$
- The probability distribution about the *current state* of the environment, the human and the robot assuming to be in s, being partially observable through observation ω. Since it is not possible to define the current state with complete certainty a belief distribution is maintained to express the history of the robotic actions and state transitions of the domain such as at time t, the robot, the human and the environment are at state s considering the sequence of past combination of actions and observations as follows:
$$b_t(s) = P(s_t = s | \omega_t, a_{t-1}, \omega_{t-1}, ..., a_0, b_0) \tag{4}$$

The explicit definition of the aforementioned probabilities indicate the design of a well-formed POMDP model, the solution of which can be achieved through the existing solvers [16]. The outcome of this solution is an action selection policy π that maximizes the sum of the expected future reward up to specific time. This policy comprises a mapping from the current state belief probability to the action space A. Given the computed policy, the robot can select an optimal action by computing its belief state based on the following update rule:

$$b'(s') = \frac{O(s', a, \omega) \sum_{s \in S} T(s, a, s') b(s)}{P(\omega | a, b)} \tag{5}$$

where b' is the updated belief, b is the given belief at the previous time step and (a, ω) is the latest combination of robot action and observation.

3.2 Design Methodology of POMDP models

Following the aforementioned theory it is revealed that the precise algorithms required for the computation of optimal policies are defined by an exponential computational growth. A single step of value iteration to compute the next selected action is on the order of $|C_t| = O(|A||C_{t-1}|^{|\Omega|})$, where $|C_{t-1}|$ corresponds to the number of components required to represent the next selected action at iteration $t-1$, while the computational burden is estimated by taking into consideration the number of iterations in each step for the $O(|S|^2 A||C_{t-1}|^{|\Omega|})$. This exponential growth for the computation of the optimal policy constrains the experts to design POMDP models limited to solve specific robotic tasks, since the number of states and actions grows drastically when trying to model real life applications scenarios by considering an abundance of environment, human and robot states, and many robotic actions that needs to be determined.

In this scope, the proposed work aims to introduce a POMDP designing methodology suitable for the decision making of multipurpose, social assistive service robots, that will be capable of resolving multiple assistive tasks, as derived in our specific case from computer vision based human activity monitoring.

This is achieved by abstracting the state and action space given the awareness of the robot for the user. Specifically, since the state space is partially observable, it can only be conceptually grouped by defining scalable blocks of states that correspond to distinct levels of robot alert $\overline{S} = \{S_H, S_M, S_L\}$. Herein, the state space is conceptually partitioned in three levels of robot alert namely *High*, *Medium* and *Low*. The states that may belong to the S_H level of robot alert group correspond to phases in the assistive task that the human requires drastic assistance from the robot. The S_M levels of robot alert define the group of states within the task, in which the robot has already been engaged in an assistive task and the levels of awareness about the human have been moderated. Last, the S_L levels of robot alert outlines these states where the assistive scenario has been resolved, the required intervention is diminished and the robot is complacent about the status of the human.

The additive value of the conceptual partitioning of the state space is that it indirectly defines groups of robotic actions, the context of which is related to the type of robot intervention required for the scenario denouement, given the current robot awareness about the human. Towards this direction, the action space is partitioned as follows $\overline{A} = \{A_T, A_C, A_M\}$. The A_T set corresponds to highly interventional robotic actions necessitated when the environment and the human is at the S_H; the A_C set reflects more discreet robotic actions when the status of the domain is assessed to be at S_M and the A_M consists of rather passive robotic actions, in essence applied when the levels of robot alert about the human are diminished, i.e. S_L. More precisely, the A_T set involves all the robotic actions required to fulfil a robot engagement to resolve a specific task i.e. navigation, manipulation, grasping, hand over, which is orchestrated by a task-specific planner. The A_C set of actions is less invasive than the A_T set and comprises the bidirectional communication planning required for the communication with the user supporting modalities such as dialogues, user interface displays, gestures

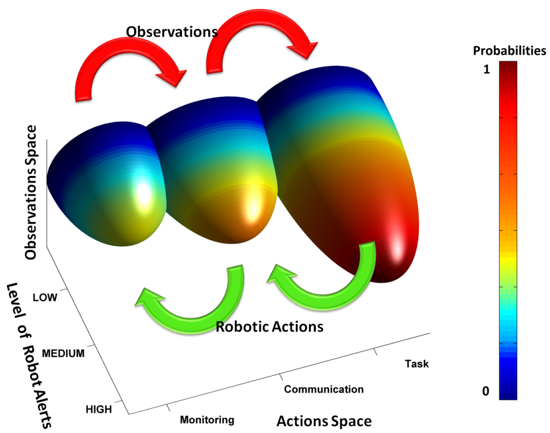

Fig. 1. The conceptual representation of the design methodology. The observations incline to transit the system to higher level of robot alerts while the decided robotic actions tend to switch the system to lower level of robot alert.

and even notification with augmented reality. The A_M set of actions corresponds to the monitoring components of the robot triggering functionalities suitable for assessment of the current status of the human and the environment, such as human detection and tracking, human cognitive and physical abilities assessment, activities interpretation and objects detection and recognition. It is revealed that this set of actions is passive, since the robot monitors the human and the environment, while the observations acquired from these actions is expected to alter the state of the domain. Although it is evident that complete situation awareness is not feasible since the environment and the human should be modeled and monitored with limited robot sensors and the acquired sensor observations are noisy, the designed POMDP model aims to propagate the system to the S_L set of states by selecting the corresponding set of actions. This feature is regulated in the proposed design methodology by carefully assigning the values at the reward function, endorsing thus the system respectively. In particular, a positive reward value is passed to the model when the selected action transits the system from higher to lower level of robot alert state, while a negative reward value is passed to the model when the selected action tends to bring the system to a higher level of robot alert. A uniform distribution is applied in the rewards function when the system passes from medium to medium levels of robot alert states. Through this methodology, the POMDP model is designed in a human-centric manner, where the partial observable set of states correspond to the status of the human and the environment, while the set of actions are solely robotic related resulting thus, in a prompting system aiming to draw decisions about the robot intervention in order to reduce the awareness of the robot about the human and thus, solve the assisting scenario.

The principal behind the design of such a POMDP model is outlined in Fig. 1, where a conceptual representation of the levels of robot alert along with the actions space and the observations space is graphically illustrated. Specifically, the figure represents all possible triplets (s, a, ω) of actions, states, and observations with respect to their probability of occurrence. The actions space is divided into three major categories which subsequently effect the conceptual clustering of the states of the domain in terms of robot alert levels. On the contrary, observations remain the free parameter that dictates the probability of simultaneous occurrence for each triplet. Additionally, the observations probabilities of occurrence are graphically illustrated in terms of a color map in Fig. 1. To this end, the immanent trend of the observations is to increase the robot's alert level i.e. transition to a state that belongs to high level of alert, while the system through actions attempts to stabilize the system in states that belong to a group of lower level of robot alert.

4 An Exemplified Case Study

At this point, it should be stressed that the work presented herein aims at defining a POMDP model-design framework, suitable for decision making in assistive service robotic applications, while the exact solver used for the policy π is out the scope of this work, as we assume that any solver is adequate to converge in a solution given a well structured POMDP model [16]. Since, to the best of our knowledge no similar work that focuses on the design methodology of POMDP models for robotic applications exists and the current applications recall the superior knowledge of the experts for such a design, a direct comparison of our method is not feasible. However, we append herein a realistic modelling of an assistive scenario utilizing our POMDP design methodology, while the validity of the derived decision making policy has been assessed with simulation and proved adequate to resolve the scenario successfully.

A challenging objective for the contemporary service robots living with elderly is to monitor their daily activities, interpret hazardous situations and to notify a relevant or ask for external help. More specifically, we consider the scenario of robot assisting elderly people during cooking activities where it is common phenomenon to forget electric appliances turned on. Specifically, during cooking activity the robot observes the human gathering the products required to prepare a specific meal. The robot monitors this activity and in case that some objects-materials are missing it asks the human if it should fetch those objects and if in cast that it receive an affirmative response it acts accordingly. In the normal situations where all objects have been successfully fetched the robot proactively examines the state of their storage place i.e. fridge, cupboards and identifies their state. The robot once again communicates with the human to notify him/her about the situation and if it is necessary it is engaged in a robotic action to close a forgotten appliance. In hazardous situations the robot assesses the risk, and decides whether external help is required utilizing an external communication channel. It is apparent that the aforementioned scenario is very complex and requires advanced decision making from the robot in order to decide when and how to intervene in order to assist the human.

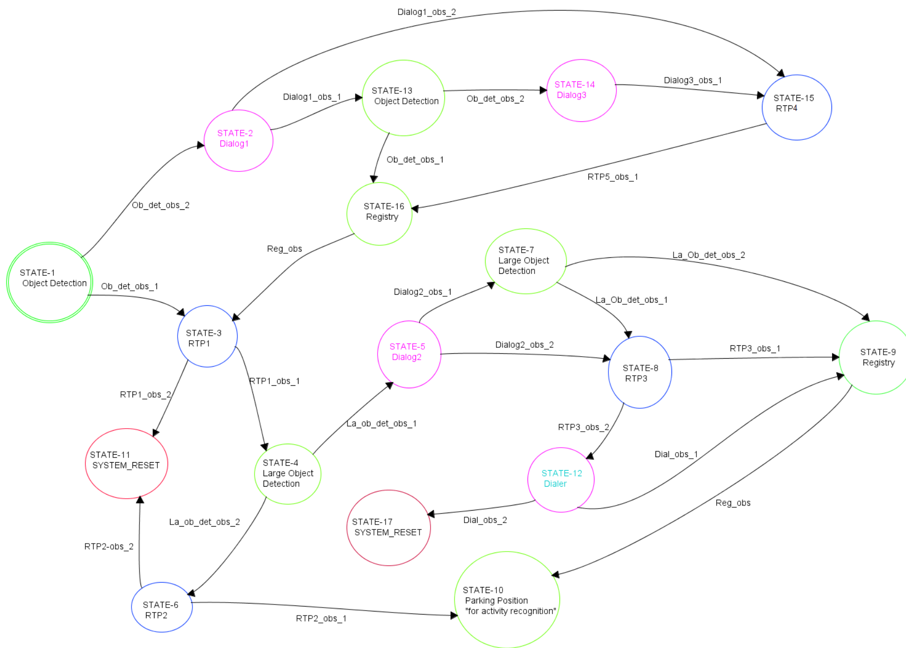

Fig. 2. An indicative state diagram that supports the robot decision making for assisting elderly during daily cooking activities. In this diagram, the task related robotic actions are highlighted with blue, the communication with the humans related actions are highlighted with magenta and the human and environment monitoring robotic actions are highlighted with green.

The flow of this scenario can be ideally described by a state diagram that conjugates the states of the domain with the robot actions as exhibited in Fig. 2, the aim of which is to identify all the required robot actions need to be modeled. In this figure, the task related robotic actions are highlighted with blue and retain the abbreviation "$RTP\#$". The robotic actions related to the direct communication with the human are highlighted with magenta and retain the abbreviations "*Dialog#*" and "*Dialer#*", while the monitoring robotic actions connected with the detection of the human and environment state are highlighted with green. A summary of the explicit interpretation of these robotic actions is provided in Table 1. Moreover, following the designed methodology described in Sec. 3.2, the identified robot actions and the domain state space are grouped accordingly. By carefully examining the diagram in Fig. 2 it is revealed that the observations acquired from the actions that belong in the set of actions A_M, tend to transit the system to higher level of alert while the observations gathered from actions linked to A_T and A_C sets, propagate the system to a lower level of robot alert. The aforementioned remark is also justified by the fact that the scenario starts and finalizes from a state that belongs to the S_L

Levels of Robot Alert			Actions		
High	*Medium*	*Low*	*Task*	*Communication*	*Monitoring*
State-3	State-2	State-1	**RTP1**: Robot navigates to the parking position suitable to monitor the state of appliance	**Dialog1**: Robot communicates with human about some missing objects and asks if it should fetch them	**Object-Detection**: The SW component suitable to detect and recognize small objects
State-6	State-5	State-4	**RTP2**: Robot navigates to the parking position suitable to monitor the cooking activity	**Dialog2**: Robot communicates with human about forgetting to turn off an appliance and asks if it should close it	**Large object detection**: The SW component suitable to recognize the state of large articulated objects
State-8	State-12	State-7	**RTP3**: Robot plans the actions for navigation and manipulation of appliance	**Dialog3**: Robot informs the human that it will go manipulate the appliance	**Registry**: The SW component suitable to register the incidents
State-15	State-14	State-9	**RTP4**: The robot fetches the missing objects	**Dialer**: The robot failed to turn off the appliance and notifies for external help	**Parking Position**: The SW component suitable to switch the robot in monitoring state where the human and environment are observed
—	—	State-10	—	—	—
—	—	State-13	—	—	—
—	—	State-16	—	—	—

Table 1. Mapping from state diagram to POMDP interpretation using Level of Robot Alerts (LoRA) and Actions. Note that $State-11$ and $State-17$ are considered to be control states and when the system prompts a respective action the robot is switched to the monitoring state while the scenario could be re-initialized.

conceptual set of domain states. The arrows that link the states are functional operators and correspond to specific observations expected to be returned after the execution of the corresponding action, which are strictly declared within the POMDP model passing increased observation probabilities to the respective (s, a, ω) triplets. During the design phase of the POMDP, such diagrams can be considered as maps that constraint the produced policy π by defining the transition probabilities among the linked states with increased probability values using the expression described in Eq.2. At the same time, the probabilities stemming from observations among linked robotic actions are also explicitly declared during the design of the POMDP passing increased values to the Eq. 3, while the rest of the observations are modeled within the POMDP as described in [19]. An important role during the design phase is play the definition of the rewards; where a very positive reward is passed to system when it transits from a state of high a state to lower level of alert, while a very negative reward in passed in the opposite situation. Finally, by using the belief state update as described in Eq. 5 the system can start from any state in the derived policy graph π, and is able to reach to a monitoring related state (S_M) due to the descending designed method of the POMDP model in terms of levels of robot alert.

5 Conclusions and Future Work

In this work, a theoretical framework for designing POMDP models suitable for multipurpose robotic applications has been introduced. Specifically, the limitation in the POMDP design due to the great amount states and robot actions that need to be modeled in real life scenarios is tackled herein with a human-centric design method where the robotic actions are decided based on the awareness of the robot about its cohabitant. The POMDP theory has been interpreted in a robotic-wise manner and the methodology presented herein is based on the conceptual abstraction of the state space using level of robot alerts, which are conjugated with respective groups of robotic actions. Through this procedure the context of the robot actions are connected to the type of assistance that is required to offer and it has been proved that intense robotic actions such as navigation and manipulation tend to transit the system at lower level of robot alerts making the robot complacent about the human. The proposed theoretic framework has been applied in a challenging scenario suitable for assistive robots by analyzing the foreseen robotic sub-tasks in a step-wise manner. Through this procedure, designing details about the parameterization of a POMDP model has been provided offering to the respective community a paradigm to design similar decision making models. Through this procedure, a native decision making system has been designed based on prompting system discharging the POMDP model from the burden to handle low-level robotic actions. In our future work we plan to extend our method by connecting the robotic high-level actions with task and communication planners with the aim to introduce a complete decision-and-act system suitable to operate on multipurpose assisting robots.

Acknowledgments

This work has been supported by the EU Horizon 2020 funded project namely: "Robotic Assistant for MCI Patients at home (RAMCIP)" under the grant agreement with no: 643433.

Bibliography

[1] Bradshaw, J.M., Feltovich, P.J., Jung, H., Kulkarni, S., Taysom, W., Uszok, A.: Dimensions of adjustable autonomy and mixed-initiative interaction. In: Agents and Computational Autonomy. Springer (2003) 17–39

[2] Goodrich, M.A., Schultz, A.C.: Human-robot interaction: a survey. Foundations and trends in human-computer interaction **1**(3) (2007) 203–275

[3] White, D.J.: A survey of applications of markov decision processes. Journal of the Operational Research Society (1993) 1073–1096

[4] Pineau, J., Thrun, S.: High-level robot behavior control using pomdps. In: AAAI-02 Workshop on Cognitive Robotics. Volume 107. (2002)

[5] Mihailidis, A., Fernie, G.R.: Context-aware assistive devices for older adults with dementia. Gerontechnology **2**(2) (2002) 173–188

[6] Hoey, J., Poupart, P., Boutilier, C., Mihailidis, A.: Pomdp models for assistive technology. In: Proc. AAAI Fall Symposium on Caring Machines: AI in Eldercare. (2005)

[7] Mihailidis, A., Boger, J.N., Craig, T., Hoey, J.: The coach prompting system to assist older adults with dementia through handwashing: An efficacy study. BMC Geriatr **8**(1) (2008) 1

[8] Hoey, J., Plötz, T., Jackson, D., Monk, A., Pham, C., Olivier, P.: Rapid specification and automated generation of prompting systems to assist people with dementia. Pervasive and Mobile Computing **7**(3) (2011) 299–318

[9] Ryu, H., Monk, A.: Interaction unit analysis: A new interaction design framework. Human–Computer Interaction **24**(4) (2009) 367–407

[10] Grześ, M., Hoey, J., Khan, S.S., Mihailidis, A., Czarnuch, S., Jackson, D., Monk, A.: Relational approach to knowledge engineering for pomdp-based assistance systems as a translation of a psychological model. International Journal of Approximate Reasoning **55**(1) (2014) 36–58

[11] Cassandra, A.R., Kaelbling, L.P., Kurien, J.A.: Acting under uncertainty: Discrete bayesian models for mobile-robot navigation. In: Intelligent Robots and Systems' 96, IROS 96, Proceedings of the 1996 IEEE/RSJ International Conference on. Volume 2., IEEE (1996) 963–972

[12] Grady, D., Moll, M., Kavraki, L.E.: Automated model approximation for robotic navigation with pomdps. In: Robotics and Automation (ICRA), 2013 IEEE International Conference on, IEEE (2013) 78–84

[13] Foka, A., Trahanias, P.: Real-time hierarchical pomdps for autonomous robot navigation. Robotics and Autonomous Systems **55**(7) (2007) 561–571

[14] Candido, S., Hutchinson, S.: Minimum uncertainty robot navigation using information-guided pomdp planning. In: Robotics and Automation (ICRA), 2011 IEEE International Conference on, IEEE (2011) 6102–6108

[15] Doumanoglou, A., Kargakos, A., Kim, T.K., Malassiotis, S.: Autonomous active recognition and unfolding of clothes using random decision forests