

Real-time hierarchical POMDPs for autonomous robot navigation

Amalia Foka*, Panos Trahanias

Institute of Computer Science, Foundation for Research and Technology – Hellas (FORTH), P.O. Box 1385, Heraklion, 711 10 Crete, Greece
Department of Computer Science, University of Crete, P.O. Box 1470, Heraklion, 714 09 Crete, Greece

Received 10 December 2004; received in revised form 8 January 2007; accepted 11 January 2007
Available online 15 February 2007

Abstract

This paper proposes a new hierarchical formulation of POMDPs for autonomous robot navigation that can be solved in real-time, and is memory efficient. It will be referred to in this paper as the Robot Navigation–Hierarchical POMDP (RN-HPOMDP). The RN-HPOMDP is utilized as a unified framework for autonomous robot navigation in dynamic environments. As such, it is used for localization, planning and local obstacle avoidance. Hence, the RN-HPOMDP decides at each time step the actions the robot should execute, without the intervention of any other external module for obstacle avoidance or localization. Our approach employs state space and action space hierarchy, and can effectively model large environments at a fine resolution. Finally, the notion of the *reference* POMDP is introduced. The latter holds all the information regarding motion and sensor uncertainty, which makes the proposed hierarchical structure memory efficient and enables fast learning. The RN-HPOMDP has been experimentally validated in real dynamic environments.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Robot navigation; Partially observable Markov decision processes (POMDP); Hierarchical POMDP

1. Introduction

The autonomous robot navigation problem has been studied thoroughly by the robotics research community over the last years. Contemporary methods for robot navigation [10, 15,7,8] do not considerably take into account the robot motion uncertainty, which may lead to the execution of false actions by the robot. Probabilistic methods that integrate uncertainty in motion planning have not been well studied until now, in contrast to probabilistic methods for mapping and localization. In this paper we introduce a Hierarchical POMDP (HPOMDP) for probabilistic navigation. Our HPOMDP formulation simultaneously addresses probabilistically all aspects of navigation, that is motion planning, localization and local obstacle avoidance.

Partially Observable Markov Decision Processes (POMDPs) provide the mathematical framework for probabilistic planning.

POMDPs model the hidden state of the robot that is not completely observable and maintain a belief distribution of the robot's state. Planning with POMDPs is performed according to the belief distribution. Therefore, actions dictated by a POMDP drive the robot to its goal but also implicitly reduce the uncertainty of its belief.

Although POMDPs successfully meet their purpose of use, they are intractable to solve with exact methods when applied to real-world environments modelled at a fine resolution. Many approximation methods for solving POMDPs are present in the literature and have also been applied to robotics problems [1, 11,19,5,20,17,18,26,4,24]. Due to the involved computational complexity, these approximation methods can only deal with problems where the size of the state space is limited to at most a few thousands states. As a result, approximation methods cannot model large real world environments at a fine resolution and hence POMDPs are mainly used as high level mission planners.

In this paper, we propose a hierarchical representation of POMDPs for autonomous robot navigation (RN-HPOMDP) that can effectively model large real world environments at a fine resolution. Moreover, the proposed RN-HPOMDP can be solved in real time. It is utilized as a unified framework for

* Corresponding address: Computational Vision and Robotics Laboratory, Institute of Computer Science, Foundation for Research and Technology, P.O. Box 1385, 711 10 Heraklion, Crete, Greece. Fax: +30 2810 391601.

E-mail addresses: foka@ics.forth.gr, foka@cs.uoi.gr (A. Foka), trahania@ics.forth.gr (P. Trahanias).

autonomous robot navigation, implying that no other external modules are used to drive the robot. RN-HPOMDP integrates the modules for localization, planning and local obstacle avoidance; it is solved online at each time step and decides the actual actions the robot performs.

In Section 2, the necessary theoretical background for POMDPs is given followed by the formulation of each element of a POMDP for the autonomous robot navigation problem in Section 3. In Section 4 the structure of the RN-HPOMDP is presented. The methodology used for learning and planning with the RN-HPOMDP is presented in Sections 5 and 6, respectively.

Two other HPOMDP approaches are currently present in the literature that are compared against the RN-HPOMDP in Section 7.1. Experimental results are presented in Section 7 and finally this paper's conclusions and future work directions are presented in Section 8.

2. Partially observable Markov decision processes (POMDPs)

POMDPs are a model of an agent interacting synchronously with its environment. The agent takes as input the state of the environment and generates as output actions, which themselves affect the state of the environment. In the POMDP framework, a system acting in the world is not guaranteed at any time to know the state of the world, i.e. which state of the environment it occupies. Hence, states are partially observable.

Formally, a POMDP is a tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \mathcal{Z}, \mathcal{O} \rangle$, where \mathcal{S} , \mathcal{A} and \mathcal{Z} are finite sets of states, actions and observations, respectively. \mathcal{T} is the *state transition function*, $\mathcal{T}(s, a, s')$, giving the probability of ending in state s' , when the agent starts in state s and takes action a . \mathcal{R} is the *reward function*, $\mathcal{R}(s, a)$, giving the expected immediate reward gained by the agent for taking an action a when it is in state s . Finally, \mathcal{O} is the *observation function*, $\mathcal{O}(s', a, z)$, giving the probability of observing z , in state s' after taking action a .

The belief state b_t of an agent is a discrete probability distribution over the set of environment states, \mathcal{S} , representing for each state the agent's belief that is currently occupying that state. The set of all possible belief states is \mathcal{B} .

A POMDP agent is composed of two components [25]: the state estimator component, that performs the belief update, and the policy component, that solves the POMDP, as explained in the following sections.

2.1. Belief update

The state estimator component of a POMDP updates the belief state of the agent every time it executes an action. Given the belief state of the agent at time t , b_t , we would like to compute the belief state at time $t + 1$, b_{t+1} , after a transition in the process where the agent occupies state s , executes an action a and perceives an observation z . The belief that the agent is in the resulting state s' is derived by:

$$b_{t+1}(s') = P(s'|z, a, b_t) = \frac{\mathcal{O}(s', a, z) \sum_{s \in \mathcal{S}} \mathcal{T}(s, a, s') b_t(s)}{P(z|a, b_t)}.$$

The denominator $P(z|a, b_t)$, is a normalizing factor and is equal to the total probability of perceiving the observation z given the previous belief state of the agent and the action it executed:

$$P(z|a, b_t) = \sum_{s' \in \mathcal{S}} \mathcal{O}(z, s', a) \mathcal{T}(s, a, s') b_t(s).$$

2.2. Solving POMDP's

Solving a POMDP amounts to computing an optimal *policy*. A *policy* is a mapping that specifies the action the agent should execute for any possible state that it might occupy. In a POMDP formulation, the true state the agent occupies is never completely known since the agent maintains a belief over all possible states. Therefore, the computed *policy* provides a mapping of belief states to actions.

The optimal action to be executed when the agent occupies a state s_t , is the one with the maximum expected accumulated reward, $E[\sum_t \gamma^t \mathcal{R}(s_t, a_t)]$, where γ is a discount factor that determines how important are the future rewards the robot will receive. If γ is zero, the robot will maximize the reward it will receive for the next time step only. The expected accumulated reward can be computed either for a specific number of steps, the *finite horizon* case, or until the agent reaches the goal state, the *infinite horizon* case.

The function that maps each state of the belief to the corresponding expected accumulated reward is called a *value function*. The t -step optimal value function [14] is constructed iteratively by *value iteration*. In the case of Markov Decision Processes (MDPs), where the agent's state is fully observable, the t -step optimal value function is formulated as:

$$V_t^*(s) = \max_{a \in \mathcal{A}} \left[\mathcal{R}(s, a) + \gamma \sum_{s' \in \mathcal{S}} \mathcal{T}(s, a, s') V_{t-1}(s') \right].$$

However, in POMDPs where the agent's state is partially observable, the value function has to be defined over the whole belief state instead of a single state as in MDPs. Hence, for POMDPs the t -step optimal value function becomes:

$$V_t^*(b) = \max_{a \in \mathcal{A}} \left[\rho(b, a) + \gamma \sum_{b' \in \mathcal{B}} \tau(b, a, b') V_{t-1}(b') \right],$$

where \mathcal{B} is the set of all possible belief states.

The transition and reward functions, $\tau(\cdot)$ and $\rho(\cdot)$ respectively, have to be defined over a belief state, b , instead of a single state, since the true state of the agent is not completely known. Hence, they are defined as: $\tau(b, a, b') = P(b'|a, b)$ and $\rho(b, a) = \sum_{s \in \mathcal{S}} b(s) R(s, a)$.

The iterative construction of the optimal value function over the set of all possible belief states \mathcal{B} is computationally an extremely expensive procedure. It has been shown that finding an exact solution of a POMDP with infinite horizon is intractable [12]. Therefore, a number of techniques have been proposed for approximating the value function. Many approximation methods are based on solving the underlying fully observable MDP [23,1,11]. More recent approximation

methods are those based on state-space compression [19], belief compression [22] and point-based value iteration where the POMDP is solved for a sampled set of belief points [4,20,24, 5].

3. Formulation of POMDPs for the autonomous robot navigation problem

In the following we present a formulation of POMDPs for autonomous robot navigation in a unified framework. The POMDP decides the actions the robot should perform to reach its goal and also robustly tracks the robot's location in a probabilistic manner. In this paper, we are interested in dynamic environments and hence the POMDP also performs obstacle avoidance. All three functionalities are carried out without the intervention of any other external module.

The elements of the POMDP, $\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \mathcal{Z}, \mathcal{O} \rangle$, are instantiated as follows:

Set of states, \mathcal{S} : Each state in \mathcal{S} corresponds to a discrete entry cell in the environment's occupancy grid map (OGM) and an orientation angle of the robot with respect to a global reference system.

Set of actions, \mathcal{A} : It consists of all possible rotation actions from 0° to 360° termed as "action angles". The discretization of the robot orientation angles and action angles depends on the number of levels of the POMDP hierarchy (see later Section 4).

Set of observations, \mathcal{Z} : The observation set is the element of the POMDP that assists in the localization of the robot, that is the belief update after an action has been taken. The set of observations is instantiated as the output of a scan matching algorithm [3,13], i.e. the robot displacement, $(dx, dy, d\theta)$, between two consecutive laser scans. Hence, observations are only robot dependent instead of environment dependent as commonly met in the literature.

Reward function, \mathcal{R} : Since the proposed POMDP is used as a unified framework for robot navigation that will provide the actual actions the robot will perform and also carry out local obstacle avoidance for moving objects, the reward function is updated at each time step. The reward function is built and updated at each time step, according to two *reward grid maps* (RGMs): a *static* and a *dynamic* [2], that determine for each state whether it is occupied by an obstacle and how far it is from the goal position.

Transition and observation functions, \mathcal{T} and \mathcal{O} : They are initially defined according to the motion model of the robot and then they are learned as explained in Section 5.

4. The robot navigation-hierarchical POMDP (RN-HPOMDP)

POMDP solution methods suffer from the "curse of dimensionality" [6] and also the "curse of history" [4].

Applying both state space and action space hierarchy, as in the RN-HPOMDP, both curses can be harnessed. In the following we present the structure of the RN-HPOMDP.

4.1. RN-HPOMDP structure

The RN-HPOMDP is built through an automated procedure using as input the map of the environment and the desired discretization of the state and action space. The map of the environment can be either a probabilistic grid map obtained at the desired discretization or a CAD map.

4.1.1. Determining the number of levels of hierarchy of the RN-HPOMDP

The RN-HPOMDP structure is built by decomposing a flat POMDP with large state and action space into multiple POMDPs with significantly smaller state and action spaces. Therefore, in levels other than the bottom level, POMDPs are composed of states and actions that have a coarse discretization and do not represent the actual state the robot occupies or the actual action the robot will perform. Hence they are termed as *abstract states* and *abstract actions* [25].

The process of building the hierarchical structure is performed in a top-down approach. The top-level of the RN-HPOMDP has a discretization of angles of 90° and at each subsequent level the discretization is doubled. Hence, the number of levels of the hierarchical structure is given by the \log_2 of the ratio of the top-level discretization and the desired discretization, ϕ , plus one level that is the top level. Thus, the number of levels of the RN-HPOMDP structure, L , will be $L = \log_2(90^\circ/\phi) + 1$.

4.1.2. Construction of the top-level of the RN-HPOMDP

The top level of the hierarchical structure is composed of a single POMDP with very coarse resolution. Hence it can represent the whole environment with a small number of abstract states. The grid resolution of the top-level states is equal to $d \times 2^{L-1}$, where d is the desired discretization of the whole RN-HPOMDP structure and L is the number of levels of the structure. The orientation angle of the robot and the action angles are also discretized in a very coarse resolution of 90° and thus represent the basic four directions $[0^\circ, 90^\circ, 180^\circ, 270^\circ]$.

The total number of states of the top-level POMDP is equal to $|\mathcal{S}^0|/2^{2(L-1)}$, where $|\mathcal{S}^0|$ is the number of states of the corresponding flat POMDP. The number of states of the top-level POMDP is reduced once by 2^{L-1} because of the coarser grid resolution and again by 2^{L-1} because of the coarser resolution of the orientation angle, as compared to the corresponding flat POMDP.

To summarize, the top level is always composed of a single POMDP with predefined discretization of the orientation and action angles at 90° . The state space size of the top-level POMDP is variable and dependent to the discretization of the corresponding flat POMDP and the number of levels of the hierarchical structure. Hence, the number of levels of the HPOMDP structure, L , should be such that it ensures that the size of the top-level POMDP remains small.

4.1.3. Construction of the intermediate levels of the RN-HPOMDP

Subsequent levels of the HPOMDP are composed of multiple POMDPs, each one representing a small area of the environment and a specific range of orientation angles. The actions of an intermediate level POMDP are a subset of the actions of the corresponding flat POMDP.

In detail, each state of the top-level POMDP corresponds to a POMDP at the immediate next level, as we go down the hierarchical structure. A POMDP at an intermediate level l , has states that represent grid locations of the environment at a resolution of $d \times 2^{(L-l)}$, where l is the current intermediate level. Thus, by going down the hierarchical structure the grid resolution of a level's POMDP is twice the resolution of the previous level.

Orientation angle decomposition

Going down the hierarchical structure, the resolution of the orientation angle is also doubled. Since the resolution of the orientation angle is increased as we go down the hierarchical structure, the whole range of possible orientation angles, $[0^\circ, 360^\circ]$, cannot be represented in every intermediate level POMDP. This would dramatically increase the size of the state space and therefore we choose to have many POMDPs that represent the same grid location but with a different range of orientation angles. The range of orientation angles that is represented within each intermediate level POMDP is expressed in terms of the orientation angle, θ_p , of the previous level state that is decomposed, and is equal to:

$$\left[\theta_p - \frac{90^\circ}{2^{l-2}}, \theta_p + \frac{90^\circ}{2^{l-2}} \right],$$

where l is the current intermediate level. By the above expression of the range of orientation angles, every intermediate level POMDP will always have five distinct orientation angles. For example, if the state of the top level POMDP, $l = 1$, has orientation angle $\theta_p = 90^\circ$, the range of orientation angles at the next level, $l = 2$, will be equal to $[0^\circ, 180^\circ]$. As mentioned earlier the angle resolution of the top level is always equal to 90° and the next level will have double resolution, i.e. 45° . Therefore, the range of orientation angles $[0^\circ, 180^\circ]$ will be represented by five distinct orientation angles. As shown in Fig. 1, the grid location represented by the top level state is decomposed into four POMDPs, where each one represents a different range of possible orientation angles. Consequently, the size of the state space for every intermediate level POMDP is constant and equal to 20, since it always has five possible orientation angles and it represents a 2×2 area of grid locations.

Action angle decomposition

Action angles are decomposed from the top level POMDP to the next intermediate level in the same manner as with the orientation angles. The resolution of the action angles at each level is the same as the resolution of the orientation angles. Hence, it is equal to $90^\circ/2^{l-1}$. As a result, a top level state is also decomposed into multiple POMDPs, each one with a different range of orientation angles but also with a different




Fig. 1. State space hierarchy decomposition. The figure depicts the decomposition of a top level state to lower level states. The top level state corresponds to four POMDPs at level 2, each one decomposing the location of the top level state into four locations, and its orientation in one of the ranges denoted by the shaded region of the circles for each POMDP. This state of decomposition continues at lower levels until the desired discretization of the environment has been reached.

range of action angles. The range of an action set is equal to

$$\left[a_p - \frac{90^\circ}{2^{l-2}}, a_p + \frac{90^\circ}{2^{l-2}} \right],$$

where a_p is the previous level action and l is the current intermediate level. The action angles set is also always composed by five distinct actions according to the above expression.

4.1.4. Construction of the bottom level of the RN-HPOMDP

The procedure described in the previous section is used to build all intermediate levels of the hierarchical structure until the bottom level is reached. Bottom level POMDPs' state and action space is discretized at the desired resolution as a flat POMDP would be discretized. The bottom level is composed of multiple POMDPs having the same properties as all other intermediate levels' POMDPs, only that the grid location the bottom level POMDPs represent is overlapping by a region r . Overlapping regions are required to be able to solve the bottom level POMDPs for border location states. Table 1 summarizes the properties of the RN-HPOMDP structure.

4.2. The reference POMDP (rPOMDP)

The RN-HPOMDP described in the previous section can cope with the computational time requirements but cannot address the memory requirements. A flat POMDP would require to hold a transition matrix of size $(|S^0|^2 \times |A^0|)$ and an observation matrix of size $(|S^0| \times |A^0| \times |Z|)$, where $|S^0|$ and $|A^0|$ are the size of the state space and action space, respectively, of the flat POMDP. The size of the observation space, $|Z|$, is the same for the flat POMDP and the RN-HPOMDP since there is no observation space hierarchy.

The RN-HPOMDP structure requires to hold the transition and observation matrices for all the POMDPs at all levels. As can be seen in Table 1, the number of POMDPs at each

Table 1
Properties of the RN-HPOMDP structure with L levels

	Top level	Intermediate level l	Bottom level
No of POMDPs	1	$ \mathcal{A}^{l-1} \times \mathcal{S}^{l-1} $	$ \mathcal{A}^{L-1} \times \mathcal{S}^{L-1} $
Size of \mathcal{S}	$ \mathcal{S}^0 /2^{2(L-1)}$	20	$5 \times (2+r)^2$
Range of orientation angles	$[0^\circ, 360^\circ]$	$[\theta_p - \frac{90^\circ}{2^{l-1}}, \theta_p + \frac{90^\circ}{2^{l-1}}]$	$[\theta_p - \frac{90^\circ}{2^{L-1}}, \theta_p + \frac{90^\circ}{2^{L-1}}]$
Resolution of orientation angles	90°	$90^\circ/2^{l-1}$	$90^\circ/2^{L-1}$
Size of \mathcal{A}	4	5	5
Range of action angles	$[0^\circ, 360^\circ]$	$[a_p - \frac{90^\circ}{2^{l-2}}, a_p + \frac{90^\circ}{2^{l-2}}]$	$[a_p - \frac{90^\circ}{2^{L-2}}, a_p + \frac{90^\circ}{2^{L-2}}]$
Resolution of action angles	90°	$90^\circ/2^{l-1}$	$90^\circ/2^{L-1}$

level is large and dependent on the size of action space and state space. Consequently, even though each POMDP’s observation and translation matrix is small, the total memory requirements would be extremely large. The RN-HPOMDP has larger memory requirements than the flat POMDP, although the flat POMDP memory requirements are already very hard to manage for large environments. For this reason, the notion of the *reference* POMDP (rPOMDP) is introduced.

The transition and observation matrices hold probabilities that carry information regarding the motion and sensor uncertainty. In the formulation of the autonomous robot navigation problem with POMDPs, as described in Section 3, transition and observation probabilities for a given action, a , and an observation, z , depend actually only on the relative position and orientation of the robot, since states and observations are robot dependent and not environment dependent.

The transition probability of a robot from a state s to a new state s' , when it has performed an action a , is only dependent on the action a . Therefore when the robot is executing an action a , the transition probability will be the same for any state s when the resulting state s' is defined relatively to the initial state s .

The probability that the robot observes a feature z , when it is in a state s and performs an action a , can also be defined in the same manner as with the transition probabilities, since the set of features \mathcal{Z} are not environment dependent. Therefore, perceived features are dependent only on the motion of the robot, i.e. the action a it performed.

The rPOMDP is built by defining a very small state space, defined as an $R \times R$ square grid (in our implementation $R = 7$) representing possible locations of the robot and all the orientation angles of the robot that would be assigned in the flat POMDP. The centre location of the state space represents the invariant state s_r of the robot. The action and observation spaces are defined in the same manner they would be defined for the corresponding flat POMDP. This rPOMDP requires to hold transition and observation matrices of size $((R \times 2^{2+L})^2 \times |\mathcal{A}|)$ and $((R \times 2^{2+L})^2 \times |\mathcal{A}| \times |\mathcal{Z}|)$, respectively. The size of the matrices is only dependent on the size of the set of actions and observations and the number of levels of hierarchy, L , since the number of levels defines the discretization of the robot’s orientation angle. By the above, it is obvious that no matter




Fig. 2. Translation and rotation of the reference POMDP transition probabilities matrix.

how big is the environment that is to be modelled with the RN-HPOMDP the use of the rPOMDP allows to have reasonably sized matrices, depending on the choice made for R , that are easy to maintain and learn.

Given the rPOMDP, transition and observation probabilities for each POMDP in the RN-HPOMDP hierarchical structure are obtained by translating and rotating the reference transition and observation probability distributions over the current POMDP state space, as shown in Fig. 2. The transfer of probabilities is performed on-line while a POMDP is solved or the robot’s belief is updated.

The transition probability for any POMDP of the hierarchical structure, $T(s, s', a)$, is equivalent to the transition probability of the rPOMDP, $T_r(s_r, s'_r, a_r)$. The reference result state, s'_r , is determined by the following equation:

$$\begin{bmatrix} x'_r \\ y'_r \\ f'_r \end{bmatrix} = \begin{bmatrix} x_r \\ y_r \\ f_r \end{bmatrix} + \begin{bmatrix} x' - x \\ y' - y \\ f' - f \end{bmatrix},$$

where, the states s, s', s_r and s'_r are decomposed to the location and orientation triplets $(x, y, f), (x', y', f'), (x_r, y_r, f_r)$ and (x'_r, y'_r, f'_r) , respectively. The reference action is determined by $a_r = a + f - f_r$.

In the same manner, the observation probability for any POMDP of the hierarchical structure, $\mathcal{O}(s, z, a)$, is equivalent to the observation probability of the rPOMDP, $\mathcal{O}_r(s_r, z_r, a_r)$.

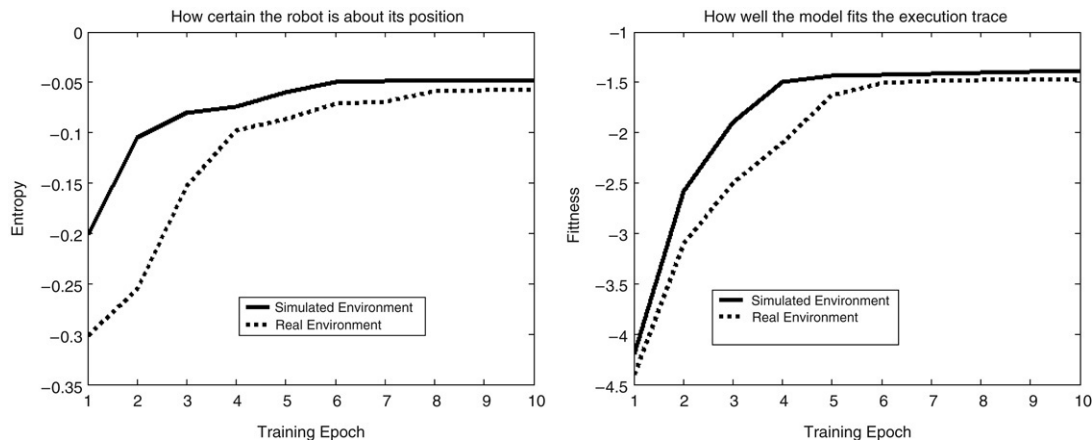


Fig. 3. Evaluation of the learned RN-HPOMDP model.

The reference observation, z'_r , is now determined as:

$$\begin{bmatrix} dx_r \\ dy_r \\ df_r \end{bmatrix} = \begin{bmatrix} d \cos(f_r + a_r) \\ d \sin(f_r + a_r) \\ df \end{bmatrix},$$

where the observations z and z_r are decomposed into (dx, dy, df) and (dx_r, dy_r, df_r) , respectively, since observations are defined as the position and angle difference between laser scans, and d is the distance $d = \sqrt{dx^2 + dy^2}$.

5. RN-HPOMDP learning

In our proposed HPOMDP structure, learning is performed only for the *reference* POMDP, since the latter transfers its learned parameters to the whole hierarchical structure. Consequently, learning is performed very fast since the rPOMDP has a very small state space. Learning the rPOMDP parameters is performed by initializing the probability matrices and adjusting their parameters iteratively according to an execution trace, that is composed of action and observation pairs, to maximize the likelihood that the execution trace was obtained by the model. The Baum-Welch [9] algorithm is utilized for this purpose.

5.1. Evaluation of the learned model

In order to test the validity of the learning procedure, we have set up an experiment aiming at a quantitative evaluation of the model that results from a learning session in specific environments. Two learning sessions have been performed; a learning session in a simulated environment where the ground truth is available and also one in a real environment. The environment chosen for both experiments is the FORTH main entrance hall, as shown in Fig. 4.

In both experiments, execution traces have been collected where the robot goes from a start state to a goal state. The start and goal states were different for each execution trace. The RN-HPOMDP for both experiments was built by discretizing the environment into 5 cm^2 cells with five levels of hierarchy, that results to a discretization step of the orientation and action angles of 5.625° . The model ‘‘appropriateness’’ has been

evaluated using the fitness and entropy measures defined in [9] as:

$$\begin{aligned} \text{fitness} &= 1/T \times \ln p(o_{1..T} | a_{1..T}) \\ \text{entropy} &= 1/(T \ln |S|) \times \sum_{t=1..T} \sum_{s \in S} [\alpha_t(s) \ln(\alpha_t(s))]. \end{aligned}$$

Fitness and entropy are indicative measures of how well the model explains an execution trace and how certain the robot is about its position. The Baum-Welch algorithm is repeated for a number of epochs until it converges. The fitness and entropy measures are graphically shown in Fig. 3 for each training epoch. Ideally, fitness and entropy should converge to zero after a sufficiently large number of training epochs. As expected, convergence to zero is not achieved, as its the case with all learning procedures. Still, after a rather small number of epochs, fitness and entropy converge to low values, indicating the validity of the learned model.

In order to provide additional quantitative results on the model accuracy, the position and orientation accuracy in maintaining the robot’s state was measured and is shown in Table 2. The peak of the POMDP’s belief distribution was used as the model’s estimate of the robot’s current state. As can be observed, the figures indicate increased accuracy of the learned model.

In the simulated environment experiments, where the ground truth is available, the position and orientation errors were measured at each time step during execution between start and goal points.

In the real environment experiments, two distinct robot locations were manually marked on the floor of the FORTH main entrance hall, as shown in Fig. 4. The robot was driven manually, as accurately as possible, to one of the marked locations and the other marked location was set as the goal position the robot had to reach. The error in the x , y location and orientation between the robot’s position after executing the trace obtained by the POMDP model and the marked location it had to reach, was measured manually as accurately as possible.

The mean position and orientation error for both experiments is very close to the discretization of the POMDP, as indicated by the entropy and fitness measures of the learned models. Both experiments, validated that the learned POMDP models




Fig. 4. The marked locations in the environment where the experimental evaluation of the RN-HPOMDP model was performed.

Table 2
Position and orientation accuracy of the learned model

Mean error	Real environment	Simulated environment
x (m)	0.053	0.023
y (m)	0.061	0.041
f (deg)	5.525	5.041

were consistent during execution in terms of maintaining the robot's belief and also in reaching the goal position.

6. RN-HPOMDP planning

Solving the RN-HPOMDP to obtain the action the robot should perform, involves solving a POMDP at each level. The intuition of the RN-HPOMDP solution is to obtain at first a coarse path that the robot should follow to reach a goal position, and then refine this path at each subsequent level in the area that the robot's current position lies, as shown in Fig. 5. The algorithm that implements the above is presented in Table 3 and its details are explained in the following.

During the RN-HPOMDP planning procedure the belief distribution of the corresponding flat POMDP is maintained at all times. This distribution will be denoted as the *full belief*. Before solving any POMDP at any level, the *full belief* is compressed, by the functions `compressTopBelief()` and `compressBelief()`, to obtain the belief distribution of the POMDP to be solved. Belief compression is performed according to the state abstraction present at each level of the RN-HPOMDP structure, i.e. the discretization reduction of each level as compared to the discretization of the corresponding flat POMDP. Therefore, the belief assigned to an abstract state, a state with coarse discretization at any level of the hierarchical structure other than the bottom level, will correspond to the average belief of all the corresponding flat POMDP states that the named abstract state has integrated. The belief distribution obtained for any POMDP is normalized before solving it.

The top level POMDP is solved, by the function `solveTopLevel()`, at an infinite horizon, until the goal state is reached. The top level POMDP produces abstract actions, i.e. actions at a coarse resolution that infer only the general direction the robot should follow and not the actual action it will perform. The abstract action to be executed, a_p , as dictated by the top level POMDP solution, determines which POMDP at the immediate next level of the hierarchical structure will be solved to obtain a new refined abstract action, that has a finer

Table 3
RN-HPOMDP planning

```

while not reached the goal state
  compressTopBelief(top level)
   $a_p = \text{solveTopLevel}(top\ level)$ 
  for  $l = 2$  to  $L$ 
    whichPOMDP = selectPOMDP( $l, a_p$ )
    compressBelief( $l, \text{whichPOMDP}$ )
     $a_p = \text{solveLevel}(l, \text{whichPOMDP})$ 
  end
  executeAction( $a_p$ )
   $z = \text{getObservation}()$ 
   $\text{belief}_L = \text{updateBelief}(\text{whichPOMDP}, a_p, z)$ 
  full belief = updateFullBelief( $\text{belief}_L, \text{whichPOMDP}$ )
end

```

discretization but still it is not the actual action the robot will perform.

The POMDP to be solved at the next level is determined by the function `selectPOMDP()`. This function searches a level l for the POMDP among all POMDPs in that level that satisfies the following two criteria:

- The zero moment of the full belief distribution over the area that is defined by the candidate POMDP states is maximum.
- The set of actions of the candidate POMDP contains an action that has minimum distance from the previous level solution's action, a_p .

The structure of the RN-HPOMDP, as described in Section 4, ensures that when solving an intermediate level POMDP the action obtained from the previous level will be refined to a new action since the action subset range is equal to

$$\left[a_p - \frac{90^\circ}{2^{l-2}}, a_p + \frac{90^\circ}{2^{l-2}} \right].$$

Therefore the solution of an intermediate level POMDP is bounded according to the previous level solution.

The described procedure continues until the bottom level is reached where an abstract action will be refined to an actual action, that is the action the robot will perform.

When the robot executes the action obtained by the bottom level POMDP solution, an observation, z , is obtained and the belief distribution of this bottom level POMDP is updated by `updateBelief()`. Bottom level POMDPs are composed of actual states and actions, i.e. subsets of states and actions that compose the corresponding flat POMDP. Hence, updating the belief of a bottom level POMDP, belief_L , amounts to updating a specific region of the *full belief*. Therefore, the

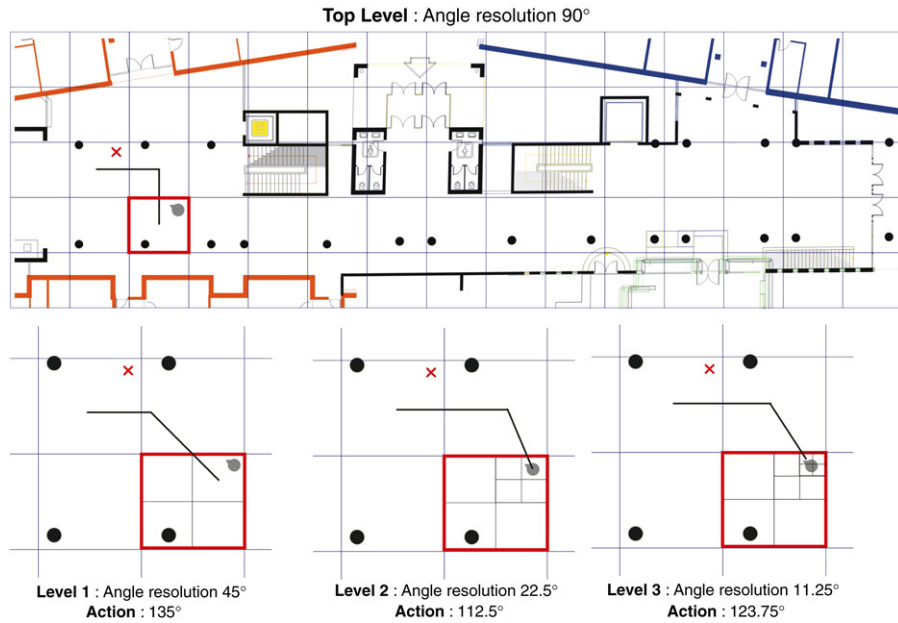


Fig. 5. Planning with the RN-HPOMDP.

belief distribution of the bottom level POMDP that was solved is transferred to the *full belief* by the function `updateFullBelief()`.

In the current implementation, all POMDPs at all levels are solved using the voting heuristic [11]. However, this is not an inherent feature of the RN-HPOMDP structure, as any other POMDP solution method can be used. Furthermore, the POMDP solution method used can also be different for each level of the hierarchical structure.

6.1. Complexity analysis

The time complexity of solving the RN-HPOMDP is obtained by determining the time complexity required to solve each POMDP at each level of the hierarchical structure.

The solution of all intermediate levels and bottom level POMDPs requires constant time, since the size of the state space and action space is always constant and predefined (cf. Table 1). Therefore, the total computational complexity required to solve the RN-HPOMDP is actually the complexity of the top level POMDP.

When approximate methods are used, such as the MLS or voting heuristic, the complexity of solving a POMDP for a single time step is $\mathcal{O}(|\mathcal{S}|^2|\mathcal{A}|)$. Hence, the corresponding time complexity of solving the RN-HPOMDP for a problem of the same size is

$$\mathcal{O}\left(\left(\frac{|\mathcal{S}|}{2^{2(L-1)}}\right)^2\right).$$

The top-level POMDP state and action space size can remain small regardless of the size of the whole environment by increasing the number of levels, L , of the hierarchical structure.

When solving a flat POMDP exactly for a single step in time t , the time complexity is $\mathcal{O}(|\mathcal{S}|^2|\mathcal{A}||V_{t-1}|^{|Z|})$, where $|V_{t-1}|$ is the number of linear components required to represent the value

function at time $t - 1$. The size of the value function at any time t is equal to $|V_t| = |\mathcal{A}||V_{t-1}|^{|Z|}$.

In contrast, the time complexity and size of the RN-HPOMDP when solved exactly for the same problem size is

$$\mathcal{O}\left(\left(\frac{|\mathcal{S}|}{2^{2(L-1)}}\right)^2 |V_{t-1}|^{|Z|}\right)$$

and $|V_t| = |V_{t-1}|^{|Z|}$, respectively.

Apart from the notable reduction in computation time due to the reduced size of the state and action space, it should be noted that the above mentioned times are for a single time step. The infinite horizon solution of a flat POMDP would require these computations to be repeated for a number N of time steps until the goal point is reached, that is dependent on the number of states of the flat POMDP, $|\mathcal{S}|$. In the RN-HPOMDP case, only the top level POMDP is solved at an infinite horizon, and the number of time steps N' until the goal point is reached, is now dependent on the number of states of the top level POMDP, $(|\mathcal{S}|/2^{2(L-1)})$.

From the above complexity analysis, we may conclude that the proposed approach takes care of the “curse of dimensionality” [6] and also the “curse of history” [4].

7. Experimental results

The RN-HPOMDP has been evaluated on its applicability as a unified model for localization, planning and obstacle avoidance in real world environments.

In terms of the performance of the RN-HPOMDP in localization, the results have been already been presented in Section 5.1, where learning of the RN-HPOMDP and its experimental validation is presented. We have presented both simulation and real-world experimental results and have shown that the RN-HPOMDP is capable of keeping track of the robot’s true position through its state estimator component.

In terms of the performance of the RN-HPOMDP in planning we present in the following a comparison with the other two hierarchical POMDP representations present in the literature. In this comparison we focus on the ability of each hierarchical representation to be applied to the problem studied in this paper, i.e. as a unified probabilistic navigation model. Furthermore, we provide a similar comparison with approximation methods for solving flat POMDPs.

7.1. Comparative results

7.1.1. Comparison with other HPOMDPs

The Theocharous HPOMDP [25] has been used as a high-level planner where the POMDP is solved once to obtain the shortest path to the goal position. In this approach the MLS heuristic is used and the time complexity for solving this HPOMDP is between $\mathcal{O}(|\mathcal{S}|^{\frac{2}{d}}N|\mathcal{A}|)$ and $\mathcal{O}(|\mathcal{S}|^2|\mathcal{A}|)$. In result, in the worst case the complexity is equal to that of solving a flat POMDP. The parameters d and N determine the complexity reduction achieved. However, these parameters depend on how well the HPOMDP was constructed and as a result this approach does not provide a guaranteed reduction.

The other HPOMDP approach present in the literature, proposed by Pineau [16], also does not have a guaranteed reduction of the action space and state space. This HPOMDP is built by grouping actions into abstract actions called subtasks. Subtasks are defined manually and according to them state abstraction is performed automatically. Therefore, this method cannot offer any bounds on the complexity times required for solving it.

The Theocharous HPOMDP has been used as high level planner using topological maps and the Pineau HPOMDP for high level robot control and dialogue management.

The state and action space reduction achieved by the above two approaches prohibits their application to the problem studied in this paper that is the unified probabilistic navigation problem using POMDPs without external modules.

7.1.2. Computational time comparison

Further to the theoretical comparison presented in the previous section, for indicative comparison purposes we provide the CPU times required to solve the RN-HPOMDP and also the Theocharous and Pineau HPOMDP approaches in Tables 4–6. It should be stressed out, that the times referring to the Pineau approach are the ones from their initial version of HPOMDP [21] where there was only action space hierarchy. It should be also noted that the CPU times mentioned are the ones the authors state and have not been obtained using computers of the same power. Another point is that the Theocharous approach is solved using the MLS heuristic and in our approach the POMDPs are solved using the voting heuristic that has the same computational complexity with the MLS heuristic. However, the Pineau HPOMDP is solved using exact methods. Regardless of the mentioned differences, the superior computational performance of our approach can be easily extracted from the tabulated results since the size of the problem is many orders of magnitude larger.

Table 4

Computation time required to solve a HPOMDP with the compared approaches

	POMDP size		CPU time (s)
	$ \mathcal{S} $	$ \mathcal{A} $	
Theocharous [25]	$ \mathcal{S} = 575$	$ \mathcal{A} = 4$	2.11–5.7
	$ \mathcal{S} = 1385$	$ \mathcal{A} = 4$	5.05–26.12
Pineau et al. [21]	$ \mathcal{S} = 11$	$ \mathcal{A} = 6$	2.84
	$ \mathcal{S} = 20$	$ \mathcal{A} = 30$	77.99

Table 5

Computation time required to solve the RN-HPOMDP with varying grid size and five levels

Grid size	POMDP size		CPU time (s)
	$ \mathcal{S} $	$ \mathcal{A} $	
5 cm × 5 cm	$ \mathcal{S} = 18,411,520$	$ \mathcal{A} = 64$	18.520
10 cm × 10 cm	$ \mathcal{S} = 4,602,880$	$ \mathcal{A} = 64$	0.911
15 cm × 15 cm	$ \mathcal{S} = 2,038,080$	$ \mathcal{A} = 64$	0.426
20 cm × 20 cm	$ \mathcal{S} = 1,150,720$	$ \mathcal{A} = 64$	0.257
25 cm × 25 cm	$ \mathcal{S} = 734,976$	$ \mathcal{A} = 64$	0.262
30 cm × 30 cm	$ \mathcal{S} = 503,808$	$ \mathcal{A} = 64$	0.251

Table 6

Computation time required to solve the RN-HPOMDP with varying number of levels and grid size of 10 cm × 10 cm

No. of levels	POMDP size		CPU time (s)
	$ \mathcal{S} $	$ \mathcal{A} $	
3	$ \mathcal{S} = 1,150,720$	$ \mathcal{A} = 16$	201.210
4	$ \mathcal{S} = 2,301,440$	$ \mathcal{A} = 32$	16.986
5	$ \mathcal{S} = 4,602,880$	$ \mathcal{A} = 64$	0.911
6	$ \mathcal{S} = 9,205,760$	$ \mathcal{A} = 128$	0.460
7	$ \mathcal{S} = 18,411,520$	$ \mathcal{A} = 256$	0.411




Fig. 6. The FORTH main entrance hall.

7.1.3. Comparison with approximation methods for solving flat POMDPs

In [5] a review of approximation methods for solving POMDPs is presented. Furthermore, one of the most recent methods for approximation is the Point Based Value Iteration (PBVI) [4] method. The time complexity of PBVI is $\mathcal{O}(|\mathcal{S}||\mathcal{A}||V_{t-1}||\mathcal{Z}||B|)$, where $|B|$ is the size of the finite set of belief points and $|V|$ remains constant throughout iterations. The time complexity of approximation methods present so far is in the best case polynomial to the size of the POMDP.

Approximation methods have been applied to problems where the POMDP comprised a few thousand states, that is as high level planners. The problem we consider consists of many orders of magnitude larger state space. As a result, the reduction of the state space that the RN-HPOMDP offers and also the reduction of the action space is crucial to its performance. Furthermore, since the proposed hierarchical structure is not restricted to a specific method for solving the underlying POMDPs, a combination of an approximation method with the



Fig. 7. Two snapshots (a)–(b) of a navigation session for avoiding a human to reach the goal position. The robot track is marked with the black squares (■) and the human track is marked with the grey dots (●).

proposed hierarchical structure can dramatically improve its performance.

7.2. Obstacle avoidance results

The RN-HPOMDP has been tested extensively in a real world environment. The robot was set to operate for more than 70 h in the FORTH main entrance hall shown in Fig. 6 in normal working hours where people were continuously present and involved with the tasks the robot was set to perform. Hence, the robot had to perform obstacle avoidance. The environment was modelled with a RN-HPOMDP of size $|S| = 18,411,520$, $|A| = 256$ and $|Z| = 24$. The RN-HPOMDP was built with seven levels. It has to be noted that the RN-HPOMDP was able to be solved in real-time with this problem size, in contrast to other hierarchical POMDPs or approximation methods that cannot accommodate large problem sizes.

In all cases the proposed navigation model has shown a robust behavior in reaching the assigned goal points and avoiding humans or other objects. Hence, the RN-HPOMDP advantages of solving globally the problems of planning and obstacle avoidance and also in a probabilistic manner can be further exploited. Since the scope of this paper is to present the RN-HPOMDP model structure how it can be solved in real time, the performance of the RN-HPOMDP as an obstacle avoider is further detailed in [2,3]. A sample path the robot followed to reach its goal and also performed local obstacle avoidance to avoid a human is shown in Fig. 7.

8. Conclusions and future work

In this work we introduced a new approach to hierarchical POMDPs (HPOMDPs). The proposed approach is designed specifically for the autonomous robot navigation problem, hence termed as Robot Navigation-HPOMDP (RN-HPOMDP). The RN-HPOMDP is utilized as a unified model that caters for planning, localization and local obstacle avoidance. Hence, it is formulated in such a manner that it does not depend on any other external modules for localization and local obstacle avoidance. To the best of our knowledge, it is the first time a POMDP has been used to provide the actual actions the robot executes and not as a high level mission planner. The

RN-HPOMDP offers significant state space and action space reduction compared to other hierarchical approaches present in the literature. Furthermore, the state space and action space reduction is guaranteed and not dependent on the environment where the robot operates. Additionally, the RN-HPOMDP can be used in conjunction with any approximation method for solving flat POMDPs, to further improve its performance. A novel approach has been also proposed for storing the model parameters with the reference POMDP (rPOMDP). The RN-HPOMDP has been experimentally validated in a real world environment.

Future work involves integrating into the RN-HPOMDP prediction about the motion of humans and other obstacles to perform efficient and effective obstacle avoidance in a predictive manner [2,3]. Furthermore, the application of the RN-HPOMDP to multi-robot navigation and cooperation will be examined.

References

- [1] A.R. Cassandra, L.P. Kaelbling, J.A. Kurien, Acting under uncertainty: Discrete bayesian models for mobile-robot navigation, in: Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems, 1996.
- [2] A. Foka, P. Trahanias, Predictive autonomous robot navigation, in: Proceedings IEEE/RSJ International Conference Intelligent Robots and Systems, IROS, 2002.
- [3] A. Foka, P. Trahanias, Predictive control of robot velocity to avoid obstacles in dynamic environments, in: Proceedings IEEE/RSJ International Conference Intelligent Robots and Systems, IROS, 2003.
- [4] G. Gordon, J. Pineau, S. Thrun, Point-based value iteration: An anytime algorithm for pomdps, in: Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI, 2003.
- [5] M. Hauskrecht, Value function approximations for partially observable Markov decision processes, *Journal of Artificial Intelligence Research* 13 (2000) 33–95.
- [6] L.P. Kaelbling, M.L. Littman, A.R. Cassandra, Planning and acting in partially observable stochastic domains, *Artificial Intelligence* 101 (1–2) (1998) 99–134.
- [7] A.A. Kassim, B.V.K. Vijaya Kumar, Path planners based on the wave expansion neural network, *Robotics and Autonomous Systems* 26 (1) (1999) 1–22.
- [8] O. Khatib, S. Quinlan, D. Williams, Robot planning and control, *Robotics and Autonomous Systems* 21 (3) (1997) 249–261.
- [9] S. Koenig, R.G. Simmons, Unsupervised learning of probabilistic models for robot navigation, in: Proceedings of the International Conference on Robotics and Automation, 1996, pp. 2301–2308.

- [10] J.-C. Latombe, Robot Motion Planning, Kluwer Academic Publishers, 1991.
- [11] M. Littman, A. Cassandra, L. Kaelbling, Learning policies for partially observable environments: Scaling up, in: Proceeding of the 12th International Conference on Machine Learning, 1995, pp. 362–370.
- [12] M.L. Littman, J. Goldsmith, M. Mundhenk, The computational complexity of probabilistic planning, Journal of Artificial Intelligence Research 9 (1998) 1–36.
- [13] F. Lu, E. Milius, Robot pose estimation in unknown environments by matching 2d range scans, Journal of Intelligent and Robotic Systems 18 (1998) 249–275.
- [14] G.E. Monahan, A survey of partially observable markov decision processes: Theory, models, and algorithms, Management Science 28 (1982) 1–16.
- [15] U. Nehmzow, C. Owen, Robot navigation in the real world: Experiments with Manchester's FortyTwo in unmodified, large environments, Robotics and Autonomous Systems 33 (4) (2000) 223–242.
- [16] J. Pineau, S. Thrun, An integrated approach to hierarchy and abstraction for POMDPs, Technical Report (CMU-RI-TR-02-21), Carnegie Mellon University, 2002.
- [17] J. Pineau, G. Gordon, S. Thrun, Applying metric-trees to belief-point POMDPs, Neural Information Processing Systems (NIPS) (2003).
- [18] J. Pineau, M. Montemerlo, M. Pollack, N. Roy, S. Thrun, Towards robotic assistants in nursing homes: Challenges and results, Robotics and Autonomous Systems 42 (3–4) (2003) 271–281.
- [19] P. Poupart, C. Boutilier, Value-directed compression of POMDPs, Neural Information Systems (NIPS) (2003).
- [20] K.-M. Poon, A fast heuristic algorithm for decision-theoretic planning, Master's Thesis, The Hong-Kong University of Science and Technology, 2001.
- [21] N. Roy, J. Pineau, S. Thrun, A hierarchical approach to pomdp planning and execution, in: Workshop on Hierarchy and Memory in Reinforcement Learning, ICML, 2001.
- [22] N. Roy, G. Gordon, Exponential family PCA for belief compression in POMDPs, Neural Information Systems (NIPS) (2003).
- [23] R. Simmons, S. Koenig, Probabilistic robot navigation in partially observable environments, in: Proceedings of the International Joint Conference on Artificial Intelligence, 1995, pp. 1080–1087.
- [24] M.T.J. Spaan, N. Vlassis, A point-based POMDP algorithm for robot planning, in: Proceedings of 2004 IEEE International Conference on Robotics and Automation, ICRA, 2004.
- [25] G. Theocharous, Hierarchical learning and planning in partially observable Markov decision processes, Ph.D. Thesis, Michigan State University, 2002.
- [26] S. Thrun, M. Beetz, M. Bennewitz, W. Burgard, A.B. Cremers, F. Dellaert, D. Fox, D. Hahnel, C. Rosenberg, N Roy, J. Schulte, D. Schulz, Probabilistic algorithms and the interactive museum tour-guide robot minerva, International Journal of Robotics Research 19 (11) (2000) 972–999.



Amalia Foka was born in Patras, Greece in 1976. She received the B.Eng. degree in Computer Systems Engineering and the M.Sc. degree in Advanced Control from UMIST, UK in 1998 and 1999 respectively. She received the Ph.D. degree in 2005 from the Department of Computer Science, University of Crete, Greece. She currently is a visiting lecturer at the Department of Computer Science, University of Ioannina, Greece. Her research interests include robotics and artificial intelligence.



Professor Panos Trahanias received his Ph.D. in Computer Science from the National Technical University of Athens, Greece. Currently he is a professor with the University of Crete, Greece and the Foundation for Research and Technology - Hellas (FORTH). From 1991 to 1993 he was with the Department of Electrical and Computer Engineering, University of Toronto, Canada, as a research associate. He has participated in many RTD programs in image analysis at the University of Toronto and has been a consultant to SPAR Aerospace Ltd., Toronto. Since 1993 he has been with the University of Crete and FORTH; currently, he is the Director of Graduate Studies, Department of Computer Science, University of Crete, and the Head of the Computational Vision and Robotics Laboratory at FORTH, where he is engaged in research and RTD projects in autonomous mobile platforms, sensory technologies, computational vision, mixed realities and human-robot interaction. Professor Trahanias has extensive experience in the execution and co-ordination of large research projects. Moreover, his work has been published extensively in scientific journals and conferences. He has participated in the programme committees of numerous international conferences; he has been the General Chair of Computer Graphics International 2004, and will be General Co-Chair of Eurographics 2008 and the European Conference on Computer Vision 2010.