

Assisting Persons with Dementia during Handwashing Using a Partially Observable Markov Decision Process

Jesse Hoey¹, Axel von Bertoldi², Pascal Poupart³, and Alex Mihailidis²

¹ School of Computing, University of Dundee jessehoey@computing.dundee.ac.uk

² Department of Occupational Science and Occupational Therapy, University of Toronto

³ David R. Cheriton School of Computer Science, University of Waterloo

Abstract. This paper presents a real-time system to assist persons with dementia during handwashing. Assistance is given in the form of verbal and/or visual prompts, or through the enlistment of a human caregiver’s help. The system uses only video inputs, and combines a Bayesian sequential estimation framework for tracking hands and towel, with a decision theoretic framework for computing policies of action – specifically a partially observable Markov decision process (POMDP). A key element of the system is the ability to estimate and adapt to user states, such as awareness, responsiveness and overall dementia level. We demonstrate the system in a set of simulation experiments, and we show examples of real-time interactions with actors.

1 Introduction

Older adults living with cognitive disabilities (such as Alzheimer’s disease or other forms of dementia) have difficulty completing activities of daily living (ADLs), and are usually assisted by a human caregiver who prompts them when necessary. The dependence on a caregiver is difficult for the patient, and can lead to feelings of anger and helplessness, particularly for private ADLs such as using the washroom. Computerized *cognitive assistive technologies (CATs)* are devices that may have the potential to allow this elderly population to complete such ADLs more independently by non-invasively monitoring the users during the task, providing guidance or assistance when necessary. This paper presents a real-time system for assisting persons with dementia during handwashing. The system was built upon three previous versions, each relaxing restrictive assumptions in previous iterations [10, 1].

Several intelligent systems that use AI techniques are currently being developed for the older adult population [6, 12]. These projects are similar to the work described in this paper in that they incorporate AI and a decision-theoretic approach. In particular, the Autominder System [13], one aspect of the Nursebot Project, applies a POMDP in the development of the planning and scheduling aspect of the system [12]. However, these systems are mainly used as scheduling and memory aids, and do not incorporate user attitude modeling or planning for prompting. Our system for assisting persons with dementia during handwashing consists of four basic components, as shown in Figure 1. Video from a camera mounted above a sink is input to a system to track objects of interest (e.g hands and towel). Object positions are passed to a *belief monitor* that estimates the progress of the user as a belief state. A *policy* then maps the belief to an action for the system to take, usually an audio or video prompt or a call for human assistance.




Fig. 1. Schematic of the system with images of test washroom.




Fig. 2. Three flocks of 5 color features, or *specks*, tracking hands and towel.

For tracking, we use a mixed-state data-driven Bayesian sequential estimation method using *flocks* of color features [3], which allow objects to be robustly tracked over long periods of time, through large changes in shape and through partial occlusions. Flocking concepts have been used to deterministically track an object with a moving camera using KLT features [7]. Our belief monitoring and policy systems use a partially observable Markov decision process (POMDP), and a heuristic for the temporal abstraction between tracking and actions. The POMDP includes a model of the user’s mental state, such as responsiveness or overall dementia level, and allows monitoring of these user attitude traits. We denote these mental states in this paper as the user’s *attitude*. Our previous work has demonstrated the handtracker [3], and a fully observable version of the POMDP model [1]. This paper makes two novel contributions. The first is a demonstration of a new POMDP model that uses only video inputs, and that can monitor (unobserved) user attitude. The second contribution is a demonstration of the complete, working system in real time with actors.

2 Hand and towel tracking

To track the hands and towel over long periods of time, we implement a particle-filter based tracker using flocks of features as our appearance model. A flock consists of a group of distinct members that are similar in appearance and that move congruently, but that can exhibit small individual differences. A flock has the properties that no member is too close to another member, and that no member is too far from the center of the flock. The flocking concept helps to enforce spatial coherence of features across an object, while having enough flexibility to adapt quickly to large shape changes and occlusions. Figure 2(a) shows three flocks of 5 color features tracking two hands and a towel. Figure 2(b)–(d) show the same three flocks later in

the sequence, during occlusions and shape changes. The flocks maintain the track, even though the object shapes have changed.

More formally, a flock, ϕ , is a tuple $\{N_f, \theta_f, \mathbf{W}, \mathbf{v}, \xi_c, \xi_u\}$ where N_f is the number of features in the flock, $\theta_f = \{\mathbf{c}_f, \Sigma_f\}$ is a *global* Gaussian color model for all flock members, \mathbf{W} is a set of N_f features, $\mathbf{w}_i = \{\mathbf{x}_i, \omega_i\}_{i=1}^{N_f}$, with image positions \mathbf{x}_i , and feature parameters ω_i that describe image appearance. We use a simple type of feature, a color *speck*, which is a set of $N_p = 4$ pixels in a 2×2 square, with a *local* Gaussian color model, $\omega_o = \{\mathbf{c}_o, \Sigma_o\}$. The likelihood of observing an image \mathbf{z} given a flock ϕ , assumes that each feature generates parts of the image independently, $L(\mathbf{z}|\phi) = \prod_{i=1}^{N_f} L(\mathbf{z}|\mathbf{w}_i, \theta_f)$. The likelihood of image \mathbf{z} , given a speck, \mathbf{w} , in a flock with color model θ_f , is a product over speck pixels of two Gaussians

$$L(\mathbf{z}|\mathbf{w}, \theta_f) \propto \prod_{j=1}^{N_p} e^{-\gamma_o \min(c_p, \frac{1}{2}(\mathbf{z}_j - \mathbf{c}_o)' \Sigma_o (\mathbf{z}_j - \mathbf{c}_o))} e^{-\gamma_c \min(c_p, \frac{1}{2}(\mathbf{z}_j - \mathbf{c}_f)' \Sigma_f (\mathbf{z}_j - \mathbf{c}_f))}$$

where $\mathbf{z}_j = \mathbf{z}(\mathbf{x}_j)$ is the image color values at speck pixel \mathbf{x}_j . The specks conform to the flock’s color model, θ_f , as well as to their local color distribution through θ_o . Finally, a constant “background” density, c_p , gives better performance under occlusions, allowing some members of the flock to be “lost” (e.g. on an occluding object). The parameters γ_o and γ_c control the tradeoff between the local and global color models. The flock’s position is updated sequentially using the standard two-step Bayesian sequential estimation recursion [2]. The dynamics of a flock is given by three terms. First, the flock members move according to some mean velocity, but with independent Gaussian noise. Second, the flock has a *collision* penalty function that varies inversely with the distance between flock members. Third, the flock has a *union* penalty function that varies proportionally to the distance between flock members and the flock mean position. The collision and union penalties are both implemented using pairwise potentials, expressed as a Gibbs distribution.

To allow for multi-modality, we use a sequential Monte-Carlo approximation (particle filter) [2], in which the target distribution is represented using a weighted set of samples. In the handwashing scenario, the tracking must be robust over long periods of time, and must be able to re-initialise if the track is lost, such as when hands leave the scene. Therefore, we augment our tracking method with mixed state dynamics [5], and a data-driven proposal [11]. A mixed-state tracker has dynamics noise, Σ_v , that varies depending on how accurately the particle filter is estimated to be tracking. The proposal uses samples generated from a combination of the dynamics process and a separate, data-driven process. Our data-driven proposal uses the feature model θ_f to build a probability map over the input image by thresholding the image, and median filtering the result to remove small components. We then choose the connected component closest to the particle being updated in this binary image and build a normalised map from which we draw flock samples.

3 POMDP model

A discrete-time POMDP consists of: a finite set S of states; a finite set A of actions; a stochastic transition model $\text{Pr} : S \times A \rightarrow \Delta(S)$, with $\text{Pr}(t|s, a)$ denoting the probability of moving from state s to t when action a is taken; a finite observation set

O ; a stochastic observation model with $\Pr(o|s)$ denoting the probability of making observation o while the system is in state s ; and a reward assigning $R(s, a, t)$ to state transition s to t induced by action a . The POMDP can be used to monitor beliefs about the system state, or to compute a *policy* that maximizes the expected discounted sum of rewards attained by the system. Since the system state is not known with certainty, a policy maps *belief states* (i.e., distributions over S) into choices of actions. We refer to [8] for an overview of POMDPs. The model we currently use is specified manually, using prior knowledge of the domain.

3.1 Handwashing States and Dynamics

The handwashing task is modeled as a POMDP with 9 state variables, 3 observation variables, and 25 actions. There are 207360 states and 198 observations. The state space can be divided into two important factors: *task* and *attitude*. We use the term *sequence* to denote a single handwashing event, and *trial* to denote a set of handwashing sequences, possibly on different days over the course of many weeks.

The *task* is described by two variables, *planstep* and *behavior*. The *plansteps* break the handwashing task down into eight situations, and the user’s *behaviors* cause transitions in the *plansteps* as shown in Figure 3. The user behaviors can be one of six activities: using soap, at water, at tap, at sink, drying, or away. Note that, whereas *planstep* is sufficient to characterise the state of the hands *behavior* is also needed to fully monitor the progress. For example, a user can be in *planstep a* with hands at the soap (trying to get to *e*) or at taps (trying to get to *b*).




Fig. 3. Simplified view of the planstep transitions for the handwashing problem. The plansteps are shown along with the state of the hands (dirty,soapy,clean,wet,dry) and the water flow (on/off). Transitions are for pairs of pre/post action behaviors for the null action. An underscore ($_$) means any behavior and $\sim b$ means any behavior other than b . A question mark, $?$ indicates a probabilistic transition.

The user *attitude* has three factors: $dementia_level = \{low, med, high\}$, giving the user’s overall level of dementia (low dementia means more demented), $awareness = \{never, no, yes\}$, telling whether the user is aware of what they are doing in the task, and $responsiveness = \{none, max, med, min\}$, giving what type of prompts the user is responsive to. We assume that $dementia_level$ does not change over a clinical trial (about 4-6 weeks). However, as we discuss in Section 3.1, and as we show in our results, our model can be used to *estimate* a particular user’s level of dementia over the course of a clinical trial. A user’s *responsiveness* is constant for a sequence, but can change from sequence to sequence (e.g. from day to day). Thus, when a user enters the washroom, we have some (fixed) prior belief about her level of responsiveness, but this prior is reset after each sequence. Finally, the user’s *awareness* can change during a sequence, depending on a number of factors, such as whether she is

given prompts, whether a human caregiver intervenes, and also her *dementia_level*. If *dementia_level* is low, the user is less likely to gain awareness. If *dementia_level* is high (less demented), the user is more likely to become and remain aware.

There are two important factors in the dynamics of the POMDP. First, the *behavior dynamics* are that the user will do the “right thing” if they are either aware, or if they are not aware, but they are responsive and have been given the correct prompt. Otherwise, they will do nothing (or something different). Second, the *planstep dynamics* are that behavior changes cause planstep transitions as shown in Figure 3. Some transitions (marked with a ? in Figure 3) include the probability that the user has abandoned her attempt. For example, if the user is in *planstep a*, and moves her hands from the taps to the sink, then the probability that she turned the water on may be less than one - she may have abandoned the task. Specific probability values can be set for particular users and tap/soap configurations.

There are three types of action the system can take: to do nothing, to prompt, or to call for human assistance. The prompts correspond to the transitions shown in Figure 3. Each prompt comes in three levels of specificity: minimal, medium, and maximal. A minimally specific prompt to use soap could be simply the verbal prompt “Use the soap now”, whereas a maximally specific version might add the user’s name, some information about location and color of the soap (e.g. “John, use the soap on your left in the pink bottle”), and may include a video demonstration.

There are three observation variables in the model. The planstep observation (PSO), is the caregiver’s indication of the planstep after an intervention. The system must gain information after an intervention to avoid repeatedly calling for assistance. The hand location observation (HL), comes from the hand tracker as described in Section 2, and gives the current locations of the hands and towel. The mean positions of the three tracks are spatially discretised into a coarse and pre-defined set of areas using threshold distances to each object (e.g. taps, soap), and combined to form the values for HL. Examples include *both_at_soap* and *water_towel* meaning one is at water and one at towel. These observations are conditioned on the *behavior*.

The temporal abstraction that maps between the video frame rate and the prompting rate is accomplished by a heuristic that updates the belief state in one of two situations. First, when the belief state is going to change significantly. Second, if the person has not changed her behavior (e.g. has not moved) for a long period of time, termed a *timeout*. These explicit timeouts are the third observation, and are an indication that the user is not aware. In the POMDP, we condition the timeout observation on the joint pre- and post-action behaviors being identical. Since the behaviors are conditioned on the awareness, a timeout will give evidence for lower awareness. Note that in some situations, the behavior may remain the same without a timeout if the hands are moving to different locations.

The POMDP model also estimates a particular user’s *attitude* over time. In particular, the model can estimate a user’s level of dementia by watching her long-term handwashing behavior over multiple sequences. The ability to estimate user traits allows the model to report such findings to carers, and can also give the model information that can be leveraged in the policy. When a new user starts using the system, the *dementia_level* variable has some prior distribution set based on the population of users. Over the course of each handwashing sequence, this distribution will shift slightly. If we then propagate this information from sequence to sequence (this is the only variable whose information is propagated), then we get a long-term estimate of the user’s dementia level. We show examples of this in Section 4.

3.2 Rewards and Solutions

The reward combines large rewards for task completion, costs for prompts proportional to specificity due to the inducement of feelings of reduced independence in the user, and large costs for caregiver calls if the user is aware or responsive.

The size of our model puts it well beyond the reach of any exact solution techniques. We used a point-based approximate solution technique based on the Perseus algorithm [15], which solves the POMDP only for a specific set of belief points. Our approach reconstructs the Perseus algorithm, taking into account the structure in the system dynamics and rewards. That is, there are many conditional independencies between variables over time, and in the reward function, that a solution technique can leverage by representing the dynamics and rewards as algebraic decision diagrams (ADDs) [14]. Our approach makes three additional approximations. First, we put a cap on the complexity of the value function (the number of alpha vectors). Typically, this bound only causes minimal decrease in the quality of a solution. Second, we merge states with values that differ by less than the Bellman error [16]. This error shrinks to zero as the computation converges, preserving optimality. Third, we only compute over observations with a significant probability of occurrence. This fast technique is related to a general method for dealing with large observation spaces [4]. We solved the POMDP using 150 alpha vectors and 65 iterations in 42 hours on a dual Intel® 2.40GHz XEON™ CPU with 4Gb of RAM, using about 2Gb of memory maximum.

We also developed a simple heuristic policy as an alternative to the computed one. The heuristic policy has a fixed set of hand-crafted thresholds on the belief distribution, and attempts to prompt when the user is not aware, and does so at whatever level of responsiveness is most likely. If the user is unaware and unresponsive, then the human caregiver is called. Other policies we compare against are the Nil policy, which always does nothing, the CG policy, which always calls the caregiver, and the certainty-equivalent (CE) policy, which looks at the most likely state given the current belief, and then acts according to the policy derived for the fully observable model (MDP). Finally, we can compare these values to the expected value achieved by the fully observable MDP with no observation noise, a utopic upper bound that may never be achievable even by an optimal POMDP policy.

4 Implementation and Results

The complete system runs on a Dell laptop with Intel® core 2 duo processor with 2GB of RAM, and uses an external monitor to play the prompts. The camera is Point Grey Research® DragonFly II™. The full system processes 640×350 frames at 19 Hz, or 320×240 frames at 40 Hz. The tracker, the belief state monitor and policy, and the prompting system operate as separate processes communicating through a UDP-based IPC mechanism. A fourth process is a central broker [9] through which all data channels pass, and that fills requests from clients for the data.

4.1 Simulations

Simulations use two models: the first is the true user model, in which we set an initial user type by specifying the user *attitude*: dementia level, responsiveness and awareness, and the second is the system’s POMDP that interacts with the true

DL	RE	AW	POMDP	Heuristic	Nil	CG	CE	fo-MDP
lo	none	never	6.4 ± 0.4	-0.7 ± 0.6	-1.8 ± 0.0	-73.3 ± 2.5	6.8 ± 0.6	9.0 ± 0.5
lo	max	no	3.1 ± 0.7	1.9 ± 1.0	-0.8 ± 0.1	-89.6 ± 3.7	3.5 ± 1.3	6.6 ± 0.6
lo	med	yes	5.6 ± 0.4	4.2 ± 0.8	0.6 ± 0.3	-115.5 ± 3.5	1.3 ± 0.8	7.8 ± 0.8
med	max	no	2.8 ± 0.5	1.3 ± 0.9	0.1 ± 0.3	-91.6 ± 3.8	3.7 ± 0.6	6.1 ± 0.6
med	min	yes	5.7 ± 1.5	6.5 ± 1.1	3.1 ± 1.8	-115.9 ± 3.4	1.7 ± 1.5	8.2 ± 0.6
hi	med	no	7.6 ± 0.6	6.3 ± 0.6	0.5 ± 0.5	-93.7 ± 3.3	7.3 ± 1.0	9.7 ± 0.5
hi	min	yes	9.8 ± 0.7	10.3 ± 0.4	10.0 ± 0.7	-116.1 ± 3.3	4.6 ± 0.5	9.3 ± 0.6
overall			6.0 ± 2.0	3.9 ± 3.6	1.0 ± 3.5	-95.3 ± 15.7	4.7 ± 3.0	8.5 ± 1.3

Table 1. Mean rewards gathered over 20 simulation trials, averaged over 10 experiments.




Fig. 4. Progression of dementia level estimates over 20 simulations for different user attitudes. Solid line: mean; dashed lines: max and min over 10 experiments.

model, and attempts to estimate the state of the user and take actions accordingly. These simulations only evaluate the decision making part of the system - they do not simulate actual video sequences or the hand-tracker behavior. For simulations regarding hand-tracking, see [3]. We evaluate the simulations by looking at the average discounted reward over time, and compare the heuristic and POMDP policies. We also look at the long term dynamics of *dementia_level*.

A simulation experiment involved a set of 20 simulations of handwashing, each for 50 steps. The dementia level belief was propagated across the 20 simulations, and we did 10 experiments with different random seeds. Table 1 shows the mean rewards averaged over the 10 experiments for representative user types, comparing the POMDP policy, the four heuristic policies, and the MDP upper bound. Table 1 also shows the average over all user types, showing that overall, the POMDP policy performs best, but not significantly better than the heuristic or certainty equivalent (CE) policy, while the call-caregiver policy (CG) is an expensive lower bound. For particular user types, we see that the CE approach does better if the user starts the trial less aware. This is because the CE approach uses a more aggressive prompting strategy due to the collapse of the belief to a single state. For example, if the belief state is close to uniform, the optimal (POMDP) policy may be to wait and see what the user does, to try to gain some information, whereas the CE approach will commit to some state, possibly causing a prompt to be issued. The CE strategy works poorly when the user is more aware. The Nil policy works best when the user is least demented and most aware (since doing nothing is close to optimal anyways).

Figure 4 shows the progression of the belief that the *dementia_level* is equal to the true dementia level over the 20 simulated sequences, averaged over 10 experiments. The maximum, mean and minimum values at each time step are shown. We see that for the extreme dementia/responsiveness/awareness levels, the POMDP learns the correct dementia level quickly. However, for intermediate dementia levels, the

POMDP learns more slowly. This is reasonable since for these intermediate cases, behaviors that could be seen in either extreme might be observed.

4.2 Actor Trials

Real-time trials were conducted in a laboratory with actors behaving according to different subject types. The camera was mounted 1.65m above the sink, and 320x180 images were processed at 47 fps. The maximum specificity prompts included video demonstrations. Four different scenarios were tested by two different actors.

In the first scenario, the user just sits and does nothing at all with her hands on the edge of the sink. The system prompts her to use the soap with medium specificity after 7 seconds, and then again with maximum specificity after 30 seconds, finally calling for human assistance after 55 seconds. The POMDP's belief in the user's *attitude* after this first episode had shifted towards *dementia_level = lo*, *awareness = never* and *responsiveness = none*. In second scenario, the user completed all steps of the task in about 35 seconds without needing any assistance. The POMDP successfully tracked her behaviors and correctly inferred that *planstep=k* was reached. The POMDP's belief in the user's *attitude* after the trial was nearly uniform over *dementia_level* (shifted from *lo* initially), and had shifted towards *aware = yes*. The belief about responsiveness did not change since no prompts were given.

Figure 5 shows the third scenario, in which the subject completes some steps without assistance, but gets stuck and requires prompting to complete others. The subject initially turns the tap on and gets her hands wet (*planstep b* at 6s) without assistance, but then does not progress to the next step. A timeout occurs at 21s, and the *awareness* begins to decrease. The subject still has not used the soap at 29s and the system prompts her to do so with medium specificity. She responds at 47s after yet another timeout, and the *responsiveness* shifts toward *max*. At 49s the user has not left the soap, so system prompts to use the pump with minimum specificity. Another prompt to turn the water off is enough for the subject to finish the task.

Figure 6 shows the fourth scenario, in which the subject requires prompting for every step and is responsive only at maximum specificity. The user does not respond to a medium specificity prompt, and the system switches to maximum specificity, which works for the remainder of the task. Notice how, between 100-120s, the system is uncertain about the planstep, due to the hands momentarily moving out of the water region. This uncertainty is resolved by 130s and the end of the task is detected.

5 Conclusions and Future Work

We have presented a system for assisting a person with dementia complete the task of handwashing that combines a flexible object tracker with monitoring and decision making using a partially observable Markov decision process (POMDP). We demonstrated the system in simulations and with actors. The system will be used in clinical trials in Toronto, Canada, in spring 2007.

The reward function, currently specified by hand, encapsulates a great deal of prior information from carers and users, that should be carefully elicited from the target population. The benefit of framing the problem using a decision theoretic model is that it provides a theoretically well founded model within which we can start to investigate questions of preference and value tradeoffs that are inherent in cognitive assistive technologies. Another benefit of the system we have developed

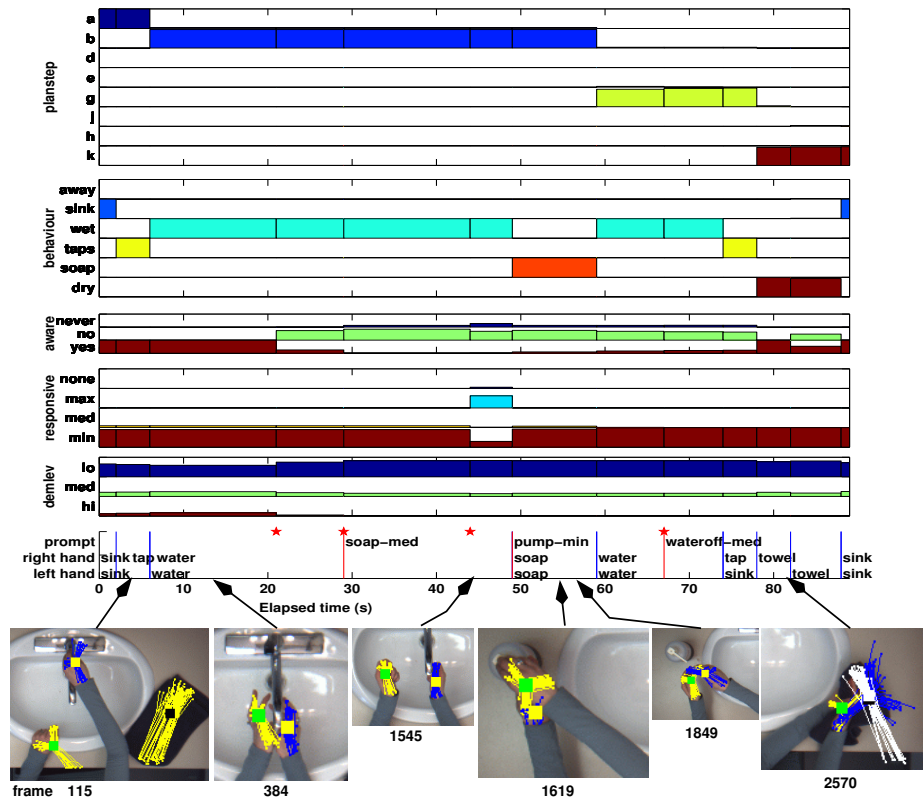


Fig. 5. Scenario three, summarized belief state, observations, timeouts (stars) and prompts. Cropped images show samples drawn from the three particle filters for hands and towel.

is its ability to generalise to other ADL. We are currently looking at implementing the same system for other important washroom ADL, such as toothbrushing, and eventually, toileting. Finally, we are investigating methods for learning the model from data, and for integrating the tracker uncertainty into the decision making.

Acknowledgments. Research supported by the American Alzheimer’s Assoc. and Intel Corp. We thank Craig Boutilier, Jen Boger, Kate Fenton, Zhonghai Li and Rozanne Wilson.

References

1. J. Boger, P. Poupart, J. Hoey, C. Boutilier, G. Fernie, and A. Mihailidis. A decision-theoretic approach to task assistance for persons with dementia. In *Proc. IJCAI*, 2005.
2. Arnaud Doucet, Nando de Freitas, and Neil Gordon, editors. *Sequential Monte Carlo in Practice*. Springer-Verlag, 2001.
3. Jesse Hoey. Tracking using flocks of features, with application to assisted handwashing. In *Proc. of BMVC*, 2006.
4. Jesse Hoey and Pascal Poupart. Solving POMDPs with continuous or large discrete observation spaces. In *Proc. IJCAI*, 2005.
5. Michael Isard and Andrew Blake. A mixed-state condensation tracker with automatic model-switching. In *Proc 6th Int. Conf. Computer Vision*, 1998.
6. H. Kautz, L. Arnstein, G. Borriello, O. Etzioni, and D. Fox. An overview of the assisted cognition project. In *AAAI Workshop on Automation as Caregiver*, 2002.

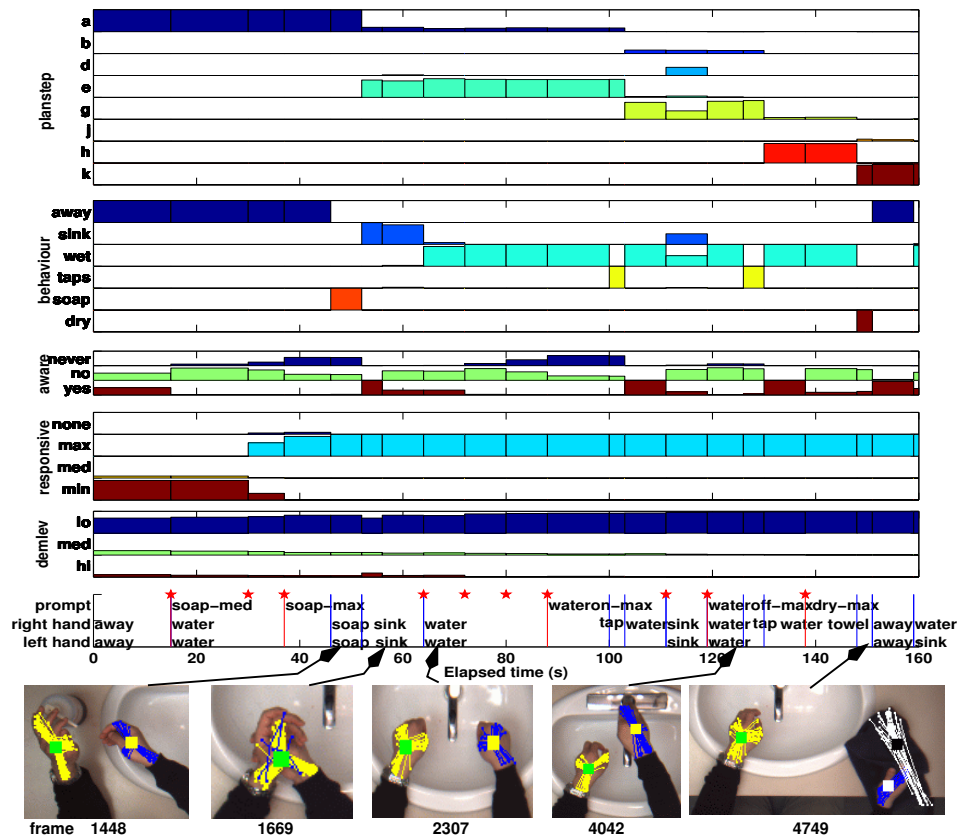


Fig. 6. Scenario three, summarized belief state, observations, timeouts (stars) and prompts.

7. Mathias Kölsch and Matthew Turk. Fast 2D hand tracking with flocks of features and multi-cue integration. In *IEEE Workshop on Real-Time Vision for HCI*, 2004.
8. W. S. Lovejoy. A survey of algorithmic methods for partially observed Markov decision processes. *Annals of Operations Research*, 28:47–66, 1991.
9. M. McNaughton, S. Verret, A. Zadorozny, and H. Zhang. Broker: An interprocess communication solution for multi-robot systems. In *Proc. IROS*, 2005.
10. Alex Mihailidis, G.R. Fernie, and J.C. Barnebel. The use of artificial intelligence in the design of an intelligent cognitive orthosis for people with dementia. *Assistive Technology*, 13:23–39, 2001.
11. Kenji Okuma, Ali Taleghani, Nando de Freitas, James J. Little, and David G. Lowe. A boosted particle filter: Multitarget detection and tracking. In *Proc. ECCV*, 2004.
12. J. Pineau, M. Montemerlo, M. Pollack, N. Roy, and S. Thrun. Towards robotic assistants in nursing homes. *Robotics and Autonomous Systems*, 42(3–4), 2003.
13. M. E. Pollack. Planning technology for intelligent cognitive orthotics. In *AIPS*, pages 322–331, Toulouse, 2002.
14. Pascal Poupart. *Exploiting Structure to Efficiently Solve Large Scale Partially Observable Markov Decision Processes*. PhD thesis, University of Toronto, 2005.
15. Matthijs T. J. Spaan and Nikos Vlassis. Perseus: Randomized point-based value iteration for POMDPs. *Journal of Artificial Intelligence Research*, 24:195–220, 2005.
16. Robert St-Aubin, Jesse Hoey, and Craig Boutilier. APRICODD: Approximate policy construction using decision diagrams. In *NIPS*, volume 14, pages 1089–1095, 2000.