Design and Prototype of a Device to Engage Cognitively Disabled Older Adults in Visual Artwork

Scott Blunsden, Brandi Richards
School of Computing
University of Dundee
Dundee, Scotland, DD1 4HN
sblunsden@gmail.com, b.h.richards@dundee.ac.uk

Tom Bartindale, Dan Jackson, Patrick Olivier
School of Computing Science
Newcastle University
Newcastle, UK, NE1 7RU
tom@bartindale.com, {d.g.jackson,p.l.olivier}@newcastle.ac.uk

Jen Boger, Alex Mihailidis
Dept. of Occupational Science and Occupational Therapy
University of Toronto, Canada, M5G 1V7
{jen.boger,alex.mihailidis}@utoronto.ca

Jesse Hoey
School of Computing
University of Dundee
Dundee, Scotland, DD1 4HN
jessehoey@computing.dundee.ac.uk

ABSTRACT

We investigate technological solutions for arts therapists who work with older adults with cognitive disabilities, such as Alzheimer's disease. We present ethnographic analysis of a survey of arts therapists in the UK and Canada, and show how there is a need for devices that can be used to promote autonomy and independence through engagement with creative visual arts. We then demonstrate a novel device that uses a touch-screen interface, and artificial intelligence software to monitor and interact with a user. Using a probabilistic model, the device monitors the behaviours of a user as well as aspects of their affective or internal state, including their responsiveness and engagement with the device. The device then uses decision theoretic reasoning to take situated actions that promote engagement from the user. We show how the device fits with the ethnographic design, and we give a laboratory demonstration of the functionality of the device. We present and discuss our next steps with this device, including end user testing.

Categories and Subject Descriptors

I.2.1 [**Artificial Intelligence**]: Applications and Expert Systems; I.2.10 [**Artificial Intelligence**]: Vision and Scene Understanding—*video analysis*

General Terms

Human Factors, Algorithms, Experimentation

Keywords

Partially Observable Markov Decision Process, MDP, POMDP, Art Therapy, Computer Vision, Face Detection, Touch Screen

PETRA'09, June 09-13 2009, Corfu, Greece

1. INTRODUCTION

This paper presents a novel tool designed to increase the capacity of art therapists to engage cognitively disabled older people in artistic activities. The tool is a touch-screen interface creative arts device that presents a user with simple creative arts tasks (e.g. painting). The device uses a camera and computer vision software to track where a person is looking, and a back end that monitors the user's activities and estimates their level of *engagement* using a partially observable Markov decision process (POMDP). The POMDP uses decision theoretic methods to reason about what actions the device can take to maintain a user's engagement. For example, the device might issue an audible prompt, or might modify the interface (e.g. by adding a new color).

The cognitive difficulties that characterize dementia include trouble following instructions, remembering steps in a process, staying engaged, and making choices. Additionally, persons with dementia often forget what they are doing and need to be reminded of their task. However, there is increasing evidence that leisurely activities promote well-being [12] and decrease dementia risk [25, 10], and that cognitive activities can slow down the progress of Alzheimer's disease [28, 4]. Engagement with visual artworks is also known to have benefits for the promotion of quality of life in older people [22]. However, many older people have difficulty motivating themselves to engage in a creative activity for a reasonable period of time. These difficulties are compounded when the older adult suffers from a cognitive disability, such as dementia (e.g. Alzheimer's disease).

Art therapists who work with older adults attempt to engage them in artistic activities, primarily with the goal of increasing their quality of life through promotion of autonomy [5] and independence, and through promotion of creative activities. Visual artwork gives older adults with dementia a way of being meaningfully engaged or occupied [13]. In turn, this leads to an increase in a person's ability to engage with their surroundings.

Art therapists at present work in day, hospital and residential care settings only. People remaining in their own home spend long periods with no occupation, as carers are often busy with daily routines. These periods reduce the ability to engage with the creative process, and can result in the person lacking motivation and desire to participate in independent activities [23]. While this engagement can be provided by a dedicated therapist, there is a lack of such therapists to support the increasing number of older adults with a cognitive disability and who are remaining in the home. Perhaps more importantly, a large benefit of engaging elderly persons with the arts at home is to enable them to do so independently and autonomously. Given the difficulties that persons with cognitive disabilities have with independent motivation and autonomy, this benefit is largely missed by persons ageing in place.

In this paper, we propose that technology can increase a therapist's ability to reach older people in their homes, providing activities that a person can engage with autonomously and independently. We demonstrate this by first discussing results of an online survey of 133 arts therapists in the United Kingdom and Canada. Applying design ethnographic methods, we translated the survey results into a set of design constraints for a device that can be used by art therapists and their clients. The results showed (a) that technological solutions would be welcomed by the arts therapy community and (b) what classes of devices would be most suitable for the target population. We have designed a prototype of such a device that fits some of these elicited design constraints, and have done basic laboratory testing to demonstrate its functionality. This device will be used to further elicit design requirements from arts therapists in targeted focus groups, and a final prototype device will be designed and tested with real users in late 2009.

This paper makes two contributions. First, we present our survey results and ethnographic analysis. Second, we present a novel prototype device that uses computer vision and decision theory to provide engaging arts activities for older adults and therapists. We present some initial testing of the device in a laboratory setting (without end users at this stage), and we discuss the next steps in the design, implementation and testing of the device.

2. RELATED WORK

The system we describe is similar to the COACH system developed for handwashing assistance [7, 16]. COACH uses an overhead camera to monitor a user by tracking their hands and the towel. A POMDP is used to estimate where the user is in the handwashing task, and audio-visual cues are delivered to assist the user in completing the task. The system we present for art therapy also uses a camera to monitor the user, but this time by watching their face. The main difference from the decision making perspective is that handwashing is a very structured task, with only few ways to accomplish it, and with goals based on physical outcomes (e.g. hands clean), as well as on user states (e.g. user independence). The creative arts task, on the other hand, is very weakly structured, with goals depending only on user internal or affective states (e.g. user engagement).

There are several other intelligent systems currently being developed for the older adult population. These include the Aware Home Project [18], the Assisted Cognition Project [11], Nursebot Project [19], the adaptive house [17], and House_n [9], 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 [20], one aspect of the Nursebot Project, applies a POMDP in the development of the planning and scheduling aspect of the system [19].

Partially observable Markov decision processes (POMDPs) provide a rich framework for planning under uncertainty[1]. In particular, POMDPs can be used to robustly optimize the course of action of complex systems despite incomplete state information due to poor or noisy sensors. For instance, in mobile robotics [19, 3] spoken-dialog systems [27] and vision-based systems for assistive technology [8], POMDPs can be used to optimize controllers that rely on the partial and noisy information provided by various sensors such as sonars, laser-range finders, video cameras and microphones.

A general POMDP model for assistive systems is presented in [6], in which the state space is broken into pieces relating to the *task*, the *user behaviour* and the *user internal states*. The same type of model is used in this paper, where *task* refers to specific interface situations, *behaviours* are a person's actions on the interface or head motions, and the *internal states* are the user's engagement and responsiveness.

3. DESIGN ETHNOGRAPHIC ANALYSIS

Ethnography is the study of human practices and interactions with the environment. Ethnographers usually study their subjects through observation, one-on-one discussions, focus groups, and surveys, and produce qualitative analyses. Design ethnographers study human interactions with objects that are situated in specific contexts, with a view to understanding the dynamic between human behaviour and the design of products and services [15]. In this Section, we report results of an application of design ethnographic methods to a survey of 133 practicing art therapists in the United Kingdom and Canada. The respondents represent a broad spectrum of specialty areas (e.g. mental illness, disabilities, cognitive impairment/dementia, etc.), specialty populations (e.g. children, adults, older adults), and preferred techniques (e.g. visual art, music, dance, etc.).

The design ethnographic analysis found three major elements of the population, and three associated design implications. These major implications are reviewed here, and are addressed by the device we propose in Section 4. Additional implications not yet addressed are reviewed in Section 6.

D1. Cognitive limitations imply Simplicity
While the goal of art therapy differs between art therapists and is often determined on a case-by-case basis, the desire to engage clients in an enjoyable and relaxing creative activity is fairly universal. The main design implication for an art therapy device is therefore that it is enjoyable to use, and that it decreases stress and anxiety. In all art therapy cases, but arguably especially for clients with dementia, it is important to avoid overwhelming and intimidating the client. This implies a design that limits options, thereby decreasing decision-making stress. However, a blank canvas can be intimidating, so an interface should always present at least a small set of choices.

D2. Physical limitations imply touch interfaces
Persons with dementia often have accompanying physical difficulties that come with ageing. Buttons and text need to be large for those with poor vision, and audio feedback and instructions need to be loud for those who are hard of hearing. Impairment of manual dexterity is a frequent obstacle for art therapists, as older clients often had trouble holding objects such as paintbrushes. Art therapists overwhelmingly agreed that a touch screen is the most viable option for use with older persons.

D3. Affective needs imply affective monitoring

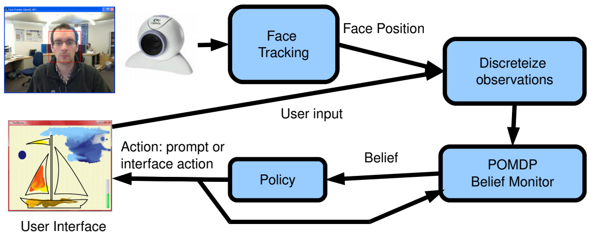

Figure 1: Overview of the system.

Art therapists actively respond to the affective signals of their clients, frequently by offering verbal encouragement if any discomfort is detected. To do this effectively, a device should be able to detect affect, either by measuring signs of discomfort or by allowing the art therapists and/or carers to input information about the client's current affective state. Furthermore, to mitigate problems of user memory, a device should react to signs of confusion, disengagement, or inactivity, by reminding clients of their task, demonstrating the task again, and/or automatically selecting options such as colors and shapes.

4. PROTOTYPE DEVICE

The prototype device aims to keep a user engaged with an artistic activity for as long as possible by monitoring their state and taking actions. This section gives an overview of the system, followed by descriptions of each component.

4.1 System Overview

A diagram of the system is given in figure 4.1. The user interface (painting program in this case) is displayed on a touchscreen monitor (design implication **D2**), specifically a NEC MultiSync 19" LCD 2010X xtra view. A web camera (Logitech Quick Cam Pro 9000) is positioned above the touch screen to view the user's face. User's interactions with the screen (position of touch), along with the presence/absence of a face, are passed to a belief monitoring system based on a partially observable Markov decision process (POMDP). A pre-computed policy is used which selects the best action to take.

4.2 User Interface: Painting

A first prototype of a user interface was designed to be simple and intuitive to use, and to provide freedom of expression for the client. The interface application presents a simple background with blobs of colored "paint". The user can blend the colors using their finger. There is always at least a background image and a few blobs of colour on the screen (design implication **D1**). The application can change its state in two ways, either by adding new paint on the screen or by changing the background image. The interface, as well as the effects of the system actions, is shown in Figure 2.

4.3 Face Detection

The device monitors a person's face, to estimate their engagement (design implication **D3**). The OpenCV [2] implementation of the Viola-Jones face detector [26] is used to detect the presence of a frontal face, as shown in Figure 3. For the purposes of this application, the detection

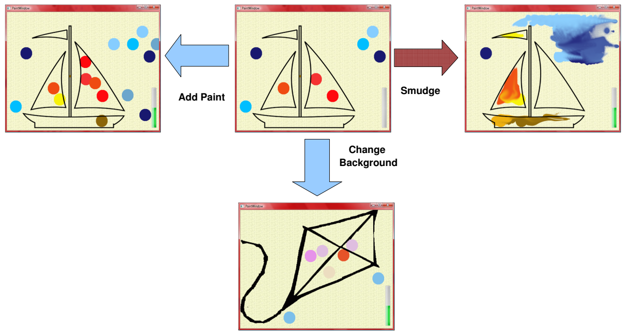

Figure 2: The application and the actions which can be performed upon it. The only user action is to smudge the colours.

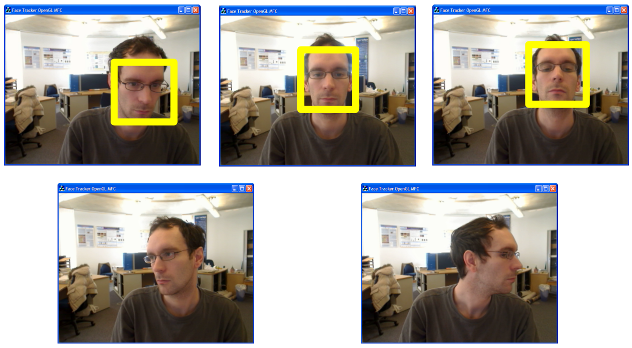

Figure 3: Face detections. Top row displays occasions when the face is successfully detected. Bottom row demonstrates when no face is detected due to out-of-plane rotation.

of a frontal face is sufficient as we only wish to detect if a person is facing the screen.

The face detection system works by computing an integral image which is an intermediate representation for the image and contains the sum of gray scale pixel values of the image. This image is used to rapidly compute rectangular Haar like features. These features sum image regions of positive and negative value. The sizes of these features is 24x24 pixels and total 160,000 in number. An AdaBoost algorithm[26] is then used to select the best features for distinguishing faces from non-faces. A cascade is used to speed up the classification by concentrating on image regions which produce features which are more likely to contain a face.

4.4 Decision Making and POMDP

A discrete-time POMDP consists of: a finite set S of states; a finite set A of actions; a stochastic transition model $\Pr : S \times A \rightarrow \Delta(S)$, with $\Pr(t|s,a)$ denoting the probability of moving from state s to t when action a is taken, and $\Delta(S)$ is a distribution over S; 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 model we currently use is specified

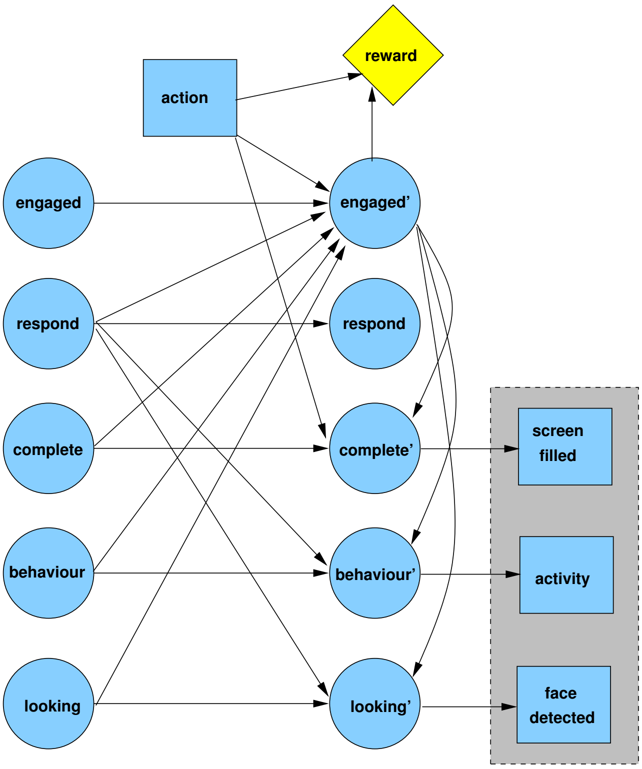

Figure 4: The POMDP model as a 2-time slice decision network: The full network can be obtained by unrolling in time. Observed variables are shown with a rectangular box. The shaded area represents the observed inputs to the model. The primed variables are those that occur after an action.

manually, using prior knowledge of the domain. However, we briefly discuss learning POMDP parameters in Section 6. We refer to [14] for an overview of POMDPs.

Intuitively, the system actions cause stochastic state transitions, with different transitions being more or less rewarding (reflecting the relative desirability of the states and the costs of taking actions). States cannot be observed exactly: instead, the stochastic observation model relates observable signals to the underlying state. The POMDP can be used to monitor beliefs about the system state using standard Bayesian tracking/filtering. Finally, a *policy* can be computed that maps *belief states* (i.e., distributions over S) into choices of actions, such that the expected discounted sum of rewards is maximized. To compute an approximate policy, we used the SymbolicPerseus package [21][1]. It implements a point-based approximate solution technique based on the Perseus algorithm [24], combined with Algebraic Decision Diagrams as the underlying data structure.

The POMDP used for modeling of user engagement in a creative artwork task is shown as a Bayesian decision network in Figure 4. We are using a factored POMDP representation in which the state space is represented as the cross product of a set of variables. The actions the system can take are to do nothing, add a blob of color, change the background and reset all colors, or to give an audio prompt.

[1] code available at www.cs.uwaterloo.ca/∼ppoupart/software

The 3 observations made by the model are

$$\textbf{activity} \in \{play, draw, intermittent, nothing\}$$
$$\textbf{screenfilled} \in \{yes, no\}$$
$$\textbf{facedetected} \in \{yes, no\}$$

activity is *play* if the user is touching the screen in a location that the system recently changed (e.g. added a color to), *draw* if the user is touching the screen in a different location, *intermittent* if the user has touched the screen in the last five seconds, but is not currently, and *nothing* if the user has not touched the screen in the last five seconds. **facedetected** is *yes* if the face tracker detects a face, and **screenfilled** is *yes* if the task is completed to a certain point (e.g. the screen is more than half filled in the painting application).

The state space contains five factors or variables. Three of these factors relate to the task:

$$\textbf{behaviour} \in \{interactive, active, inactive\}$$
$$\textbf{looking} \in \{yes, no\}$$
$$\textbf{complete} \in \{yes, no\}$$

The **behaviour**$\in\{interactive, active, inactive\}$ of a person is whether they are actively doing something on the interface, and is inferred from our observations of their finger interactions. The **looking**$\in\{yes, no\}$ variable shows if a person is looking at the screen. **behaviour** and **looking** are inferred from the corresponding observations, due to noise in the interface (touch screen errors) and in the face tracking algorithm. The **completion**$\in\{yes, no\}$ estimates whether the user has completed a certain application. This is currently inferred from the amount of screen that is filled, but could be based on a number of factors. For example, the screen could be filled with paint but the person could still be engaged in smudging paint around.

The other two factors relate to internal, affective or mental states of the user:

$$\textbf{engagement} \in \{yes, confused, no\}$$
$$\textbf{respond} \in \{cue\_color, cue\_only, color\_only, no\}$$

The user's **engagement** is the key element of this model, as it is the primary purpose of the device (to maintain engagement). A user can be engaged (*yes*, *confused*), or disengaged (*no*). We also model the user's responsiveness to the system's actions to give a prompt or cue, or to add a color (or both).

The dynamics of the POMDP hinges on the user's **engagement**: this is what is effectively changed by the system's actions, and determines the future behaviours of the user through the probability distributions over the state variables. For example,

$$P(\textbf{looking}'|\textbf{looking}, \textbf{engaged}', \textbf{respond})$$

captures that a user is more likely to look at the screen if they are engaged, or if they are not engaged, but are responsive and were just given a prompt. Similarly,

$$P(\textbf{behaviour}'|\textbf{behaviour}, \textbf{engaged}', \textbf{respond})$$

captures that a user is likely to interact with a color if they are responsive, for example. The user's engagement changes dynamically over time as a function of the system's actions, and their previous behaviours. For example, if the user is disengaged, but is looking a the screen, and the system adds a new color, then the user may become engaged with some probability. On the other hand, if a user is already engaged, and the system gives a prompt or changes the interface, the user may become confused. The user's responsiveness comes into play when they are prompted or when the interface is changed by the system. If they are responsive, the effect of the prompt is to increase their engagement. Otherwise, the system action has little or no effect.

The reward function is designed so the system maximises engagement of the user. The system is rewarded if the user is engaged (+10), and penalised if the user is confused (-1) or not engaged (-2). Motivational prompts are slightly costly (cost of 0.5), but only if the user is engaged. This models the supposed effect of a prompt reducing feelings of independence in a user if they are already engaged. Note that this cost is separate from the indirect costs incurred if a user is prompted when engaged that happen because the user may get confused, as described above.

5. EXPERIMENTS AND RESULTS

The POMDP was solved using 1000 belief points and 200 iterations, with a maximum value function size of 150 alpha vectors. The belief points were re-generated every 50 iterations (a *round*), using a policy that randomly selects between doing choosing a random action, following the equivalent MDP policy based on the most likely belief state (certainty equivalence), or following the best POMDP policy found thus far (in the first round, this is the equivalent MDP policy based on the most likely belief state, also known as the certainty equivalent policy).

The policy roughly acts according to the following strategy. In situations where the user is active, the device leaves them alone (does nothing). If the user becomes intermittently active or inactive for a period of time, the device adds some color or prompts them. If the user stops looking at the screen, the device prompts them. Finally, if the user fills the screen and becomes disengaged, the system resets.

We now present three demonstrations in the laboratory with a non-demented subject acting according to defined situations. These are meant only to illustrate the functionality of the device and user monitoring, not as end-user tests. We have not tested with end users as of yet.

5.1 Responsive person

The example in figure 5 shows a person who is responsive and engaged. The system initially plays an audio prompt giving instructions on what to do. The system observes that a person's face is present but that they haven't used the application in the last 5 seconds since the prompt was issued. Their behaviour is determined to be inactive (In). The action to add a blob of colour is then taken. The observation of the activity changes from nothing to play indicating that the person has touched the colour. The system's belief in the person's responsiveness to colors (colour respond) is increased based on this one experience. The person is also believed to be more engaged (engaged = Y) and interactively using the system (behaviour=I). The person continues, increasing the belief that the user is engaged and active. When the person is engaged the system does nothing.

5.2 Looking Away

In figure 6 we look at the system's response to a person looking away. It could also be considered a demonstration of how the face tracking will effect the systems internal beliefs.

Throughout the sequence, the system believes that the user is active (A). When no person is detected at time 2 and 3, the system estimates they are not looking with high probability. Their engagement is also estimated to be confused as they are using the system but not looking at it, indicating they may be confused as to what is going on. The system tries to repeatedly prompt them to look at the screen.

At time 4 the observation records that the person is looking at the screen and so attempts to increase engagement by giving them a visual stimulus of changing the background.

5.3 Change Background

Figure 7 displays the state of a POMDP model when a user completes a task. A colour has just been added, and it is observed that the person has interacted with it. This causes the behaviour to give a high belief for the interactive state. The person is also engaged in the activity and responsive to new colours. For the next two timesteps the person is engaged so no action is taken. For step 3 the screen is full (as given by the yes observation) and the system estimates that the activity is complete. The model selects a colour to add to the screen to see if the user touches it. Due to the user not responding to the new colour and an estimated completion of the screen the background is changed. The model does not just respond to the completion variable but attempts to keep the user engaged by providing stimulus such as a new colour or a change of background.

6. DISCUSSION

The survey analysis pointed to a number of other issues that need to be addressed in more detail. These were

D4. Saving and reviewing work

Much of the therapeutic benefit of artistic expression is the satisfaction that the artist feels upon completion of their artwork. Beyond the affective considerations for allowing clients engage with their creations, the works of art are used by therapists for a variety of reasons, including prompting of memory recall. Although persons with dementia may not remember having made their artwork, it is important that the work be saved, and for a user to see the products of their creative efforts again and again.

D5. Tangible interfaces

The use of a touchscreen interface raised additional design concerns, namely (1) how to ensure the user understood how to use the device (2) how to preserve the sensory components of the art making process. The proposed solution to the first concern is to design an interface to match the real world action of the art making activity. For example, the device could respond to the natural movement of pouring paint. To address the second concern, the device should incorporate the tactile into the activity, and should link the art materials to the sounds they would naturally make, e.g. link spray paint to a spraying sound.

D6. Art therapist involvement

The goal of an art therapy device should not be to replace the art therapist, as this relationship is a crucial component of the effectiveness of the therapy. In fact, the program could benefit from the therapist' case-by-case expertise by allowing him or her to participate in shaping the activities that are best suited to individual clients. This means designing the device in such a way that the art therapist has control over the activities the client can choose to do, includ-

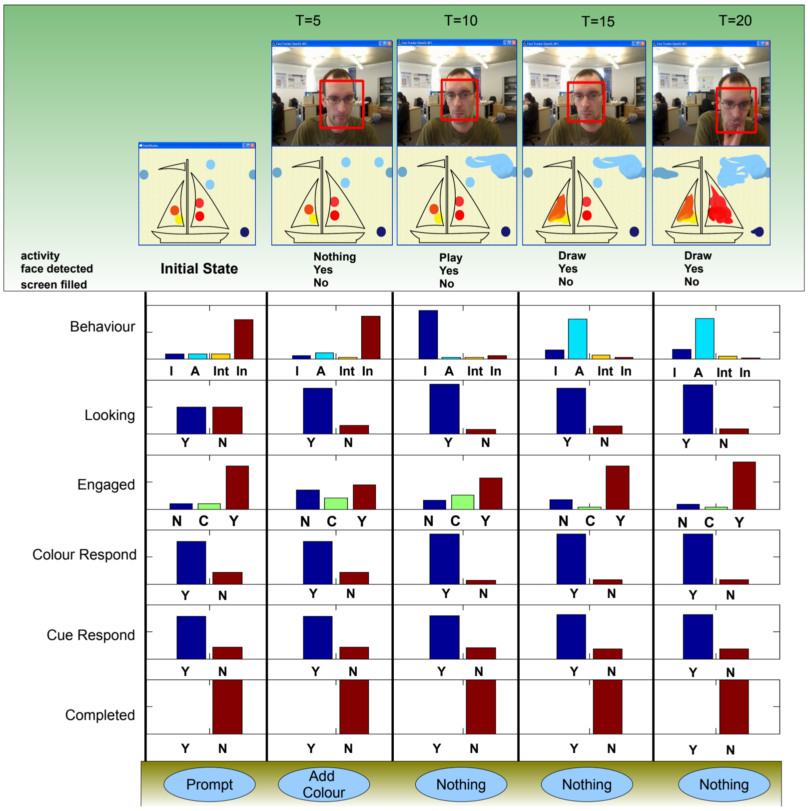

Figure 5: The state of the POMDP model. The observations are shown in the top shaded box for each timestep. The initial action is taken before any observations are made. The variable behaviour has the states interactive (I), active (A), intermittent (Int) and inactive (In). The variable engaged is modelled with the states no (N), confused (C) and yes (Y). All other variables are modelled as having a yes (Y) or no (N) state. The respond variable is split in two showing responsiveness to each type of prompt. The actions taken by the system are displayed in the bottom shaded panel. Time runs from left to right, in 5 second increments. The marginal beliefs (in $[0, 1]$) are shown for each variable as bar plots.

ing level of difficulty. Our work will address this by allowing art therapists to program the device at a high level, in order to implement their own arts tasks for their clients. The idea will be to find some underlying structure or invariants in these arts tasks that can be used as the backbone for the artificial intelligence back-end. An arts therapist will then be provided with the ability to create new arts tasks, so long as these can be mapped to this conceptual abstraction over which the decision making and user monitoring system can operate. Our plan is to approach this with targeted focus groups, followed by user testing with wizard-of-oz studies.

D7. Feedback

The device should provide feedback to the art therapists in order to inform them of their clients' progress. Therapists surveyed articulated their desires to know the following:

1. what activities were done on the computer,
2. how long the client spent doing art,
3. the time that the device was used,
4. who else used the device,
5. whether the client worked alone or with assistance, and
6. the number of prompts required.

A therapist should be able to see trends in these features over time in order to assess progress or a decline in ability. It is also clear from the survey that, compared with the final creative output, it is as important - if not more so - for the art therapist to be able to see the process of art making. This might mean incorporating a recording feature into the design, either a video recorder or a real-time recorder of touch screen activity that the art therapist can play back.

We are also integrating learning into the device. Based on a history of interactions between a client and the device, we can use standard machine learning techniques to refine the estimates of POMDP dynamics. These new estimates will then allow the device to adapt over time to its user, and the therapist to see these adaptations as additional feedback.

7. CONCLUSION

This paper has presented an analysis of a survey of art therapists that supports the claim that technological solutions are desired for promotion of engagement with creative activities in older adults with cognitive disabilities. The paper then described a novel device that uses computer vision and decision theory to monitor a user and to take actions meant to engage the user in a creative arts tasks. We showed how this device satisfied some of the constraints elicited from therapists, and discussed how it could incorporate the remainder. Our next steps are to engage with therapists,

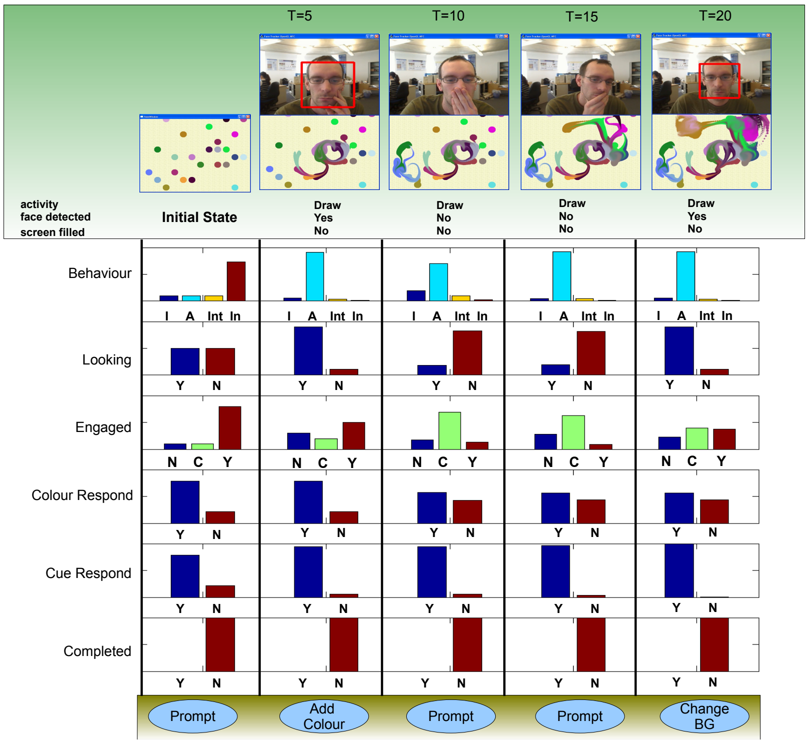

Figure 6: The state of the POMDP model when a person looks away. See Figure 5 for symbols' description.

carers and end-users more directly in a set of focus groups planned for Spring 2009, and to use the results to design final prototypes that will be tested with real end-users.

Acknowledgements
The authors thank Laurel Young, Jane Burns, and all survey participants. This work was supported by the American Alzheimer's Association under the *Everyday Technologies for Alzheimer's Care* program.

8. REFERENCES

[1] K. J. Åström. Optimal control of Markov decision processes with incomplete state estimation. *J. of Math. Anal. and Appl.*, 10:174–205, 1965.

[2] G. Bradski and A. Kaehler. *Learning OpenCV*. O'Reilly Media, 2008.

[3] T. Darrell and A. P. Pentland. Active gesture recognition using partially observable Markov decision processes. In *13th IEEE Intl. Conference on Pattern Recognition*, Vienna, Austria, 1996.

[4] Y. E. Geda, *et al.* Cognitive activities are associated with decreased risk of mild cognitive impairment: The mayo clinic population-based study of aging. In *Abstracts of the 61st Annual Meeting of the American Academy of Neurology*, Seattle, WA, 2009.

[5] J. Harlan. Beyond the patient to the person aspects of autonomous functioning in individuals with mild to moderate dementia. *The American Journal of Art Therapy*, 28(4):99–106, 1990.

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

[7] J. Hoey, P. Poupart, A. von Bertoldi, T. Craig, C. Boutilier, and A. Mihailidis. Automated handwashing assistance for persons with dementia using video and a partially observable markov decision process. *In Press: Computer Vision and Image Understanding*, 2009.

[8] J. Hoey, A. von Bertoldi, P. Poupart, and A. Mihailidis. Assisting persons with dementia during handwashing using a partially observable Markov decision process. In *Proceedings of the International Conference on Vision Systems*, Mar., 2007.

[9] S. S. Intille. The goal: smart people, not smart homes. In *Proceedings of the International Conference on Smart Homes and Health Telematics*. IOS Press, 2006.

[10] A. Karp, S. Paillard-Borg, H.-X. Wang, M. Silverstein, B. Winblad, and L. Fratiglioni. Mental, physical and social components in leisure activities equally contribute to decrease dementia risk. *Dementia and Geriatric Cognitive Disorders*, 21(2):65–73, 2006.

[11] H. Kautz, L. Arnstein, G. Borriello, O. Etzioni, and D. Fox. An overview of the assisted cognition project. In *AAAI-2002 Workshop on Automation as Caregiver: The Role of Intelligent Technology in Elder Care*, Edmonton, 2002.

[12] W. Kinney and C. Cole. Predicting life satisfaction among adults with physical disabilities. *Archives of Physical Medecine and Rehab.*, 73:863–869, 1992.

[13] T. Kitwood. *Dementia Reconsidered*. Routledge, London, 1997.

[14] W. S. Lovejoy. A survey of algorithmic methods for partially observed Markov decision processes. *Annals of Operations Research*, 28:47–66, 1991.

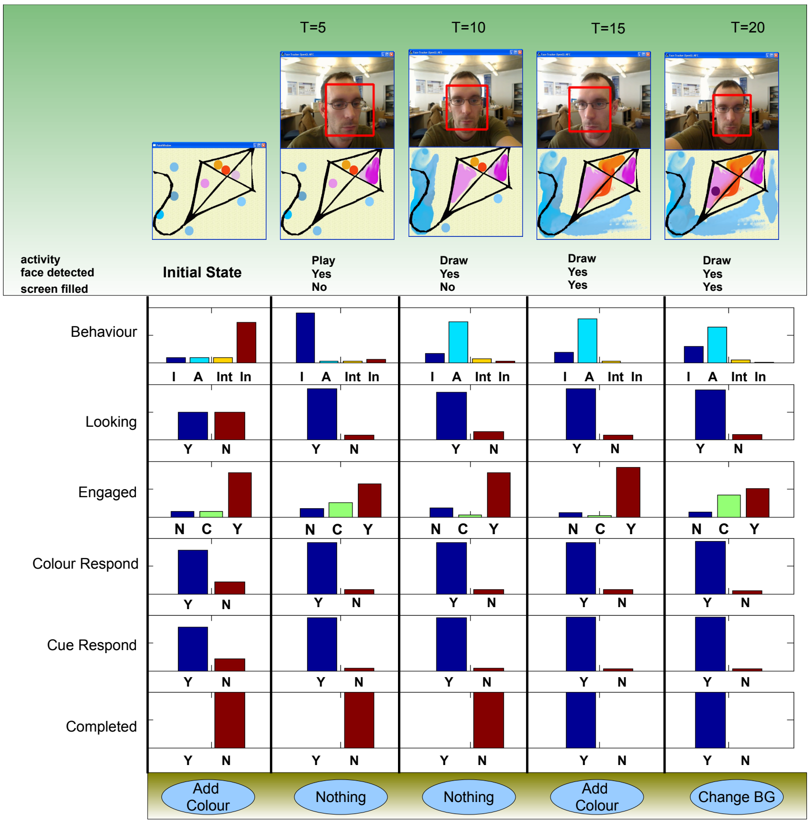

Figure 7: A person is engaged in the task until the task is completed, the background is then changed. See Figure 5 for description of symbols.

[15] C. Macaulay, D. Benyon, and C. Crerar. Ethnography, theory and design: From intuition to insight. *Intl. J. of Human-Computer Studies*, 53:35–60, 2000.

[16] A. Mihailidis, J. Boger, M. Candido, and J. Hoey. The coach prompting system to assist older adults with dementia through handwashing: An efficacy study. *BMC Geriatrics*, 8(28), 2008.

[17] M. C. Mozer. Lessons from an adaptive house. In D. Cook and R. Das, editors, *Smart environments: Technologies, protocols, and applications*, pages 273–294. J. Wiley and Sons, Hoboken, NJ, 2005.

[18] E. Mynatt, I. Essa, and W. Rogers. Increasing the opportunities for aging in place. In *Proc. of ACM Universal Usability Conference*, Arlington, VA., 2000.

[19] J. Pineau, M. Montemerlo, M. Pollack, N. Roy, and S. Thrun. Towards robotic assistants in nursing homes: Challenges and results. *Robotics and Autonomous Systems*, 42(3–4), 2003.

[20] M. E. Pollack. Autominder: A case study of assistive technology for elders with cognitive impairment. *Generations: The Journal of the American Society on Aging*, 30(2):67–69, 2006.

[21] P. Poupart. *Exploiting Structure to Efficiently Solve Large Scale Partially Observable Markov Decision Processes*. PhD thesis, U. of Toronto, Canada, 2005.

[22] J. Rusted, L. Sheppard, and D. Waller. A multi-centre randomized control group trial on the use of art therapy for older people with dementia. *Group Analysis*, 39(4):517–536, 2006.

[23] M. Seligman, editor. *Helplessness: On Depression, Development, and Death*. W.H. Freeman, San Francisco, 1975.

[24] M. T. J. Spaan and N. Vlassis. Perseus: Randomized point-based value iteration for POMDPs. *Journal of Artificial Intelligence Research*, 24:195–220, 2005.

[25] Verghese, R. B. Lipton, M. J. Katz, C. B. Hall, C. A. Derby, G. Kuslansky, A. F. Ambrose, M. Sliwinski, and H. Buschke. Leisure activities and the risk of dementia in the elderly. *New England Journal of Medecine*, 348(25):2508–2516, 2003.

[26] P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. In *In IEEE Conference on Computer Vision and Pattern Recognition*, pages 609–615, 2001.

[27] J. Williams, P. Poupart, and S. Young. Factored partially observable Markov decision processes for dialogue management. In *Proc. IJCAI Workshop on Knowledge and Reasoning in Practical Dialogue Systems*, pages 76–82, Edinburgh, Scotland, Aug. 2005.

[28] R. S. Wilson, C. F. M. de Leon, L. L. Barnes, J. A. Schneider, J. L. Bienias, D. A. Evans, and D. A. Bennett. Participation in cognitively stimulating activities and risk of incident Alzheimer disease. *Journal of the American Medical Association*, 287(6):742–748, Feb. 2002.