ePAD: Engaging Platform for Art Development

Jesse Hoey, Brandi Richards, Scott Blunsden
School of Computing,
University of Dundee, Dundee, UK

Jane Burns
School of Health Sciences,
Queen Margaret Univ., Edinburgh, UK

Tom Bartindale, Dan Jackson, Patrick Olivier
School of Computing Science
Newcastle University, Newcastle, UK

Jen Boger, Alex Mihailidis
Dept. of Occupational Therapy,
University of Toronto, Toronto, Canada

Abstract

We present a class of devices for use by art therapists working with older adults with a progressive illness such as Alzheimer's disease. We call these devices ePADs. An ePAD combines a touch-screen interface with intelligent user modeling and sensing through cameras using computer vision. Using a probabilistic model, an ePAD 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 ePAD then uses decision theoretic planning to enable situated, adaptive strategies for interaction with a human user. In this paper, we discuss results and analysis of a survey of arts therapists, and of one-on-one interviews. We then give details of the ePAD class, framed as a partially observable Markov decision process, or POMDP. A key element of this class is that instantiations can be easily made for a wide range of customisable devices and interface applications for art-making moderation. We show examples of particular instances of this model on three devices and with three different interfaces. We give laboratory demonstrations of the functionality of the devices, and we present and discuss our next steps, including end user testing.

1 Introduction

This paper presents a novel class of tools designed to increase the capacity of art therapists to engage cognitively disabled older people in artistic activities. The tools, known as *ePADs*, are touch-screen interface creative arts devices that present a user with simple creative arts tasks (e.g. painting). An ePAD uses a camera and computer vision software to track what a person is doing, and a back end that monitors the user's activities and estimates, among other things, 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). An ePAD can use its user models to adapt its strategy over time, or to aid a therapist with assessment. ePADs are also customisable by art therapists, allowing them to add new activities, or change the device's policy.

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 decrease dementia risk [Karp *et al.*, 2006], and that cognitive activities can slow down the progress of Alzheimer's disease [Wilson *et al.*, 2002; Geda *et al.*, 2009]. Engagement with visual artworks is also known to have benefits for the promotion of quality of life in older people [Rusted *et al.*, 2006]. 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 progressive illness, such as Alzheimer's disease.

Art therapy is a triadic relationship in which client, therapist and artwork engage with each other. The primary goal of art therapy with older adults is to increase quality of life through promotion of autonomy [Harlan, 1990] and independence, and through promotion of creative activities. Engaging in the art making process also provides an outlet for the person's emotional state. Through the artwork, the person can communicate, for example, his/her feelings of isolation and loneliness. In this way, the artwork is a powerful tool for connecting the person to a sense of self, others, and the environment [Burns, 2009].

Art therapists primarily work in residential, hospital and day care settings. 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 [Seligman, 1975]. While this engagement can be provided by a dedicated therapist, there is a lack of such therapists to support the increasing number of persons with a progressive illness 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 a progressive illness such as Alzheimer's disease 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, coupled with our expert knowledge of arts therapy specifically for dementia [Burns, 2009]. Applying design ethnographic methods, we uncovered a set of design constraints for devices that can be used by art therapists and their clients. We have designed a class of devices that fits some of these elicited design constraints, and have done basic laboratory testing to demonstrate the functionality of three instances of this class. Importantly, our devices are meant as more than just games for leisure: they are specifically aimed at art therapists as tools for formal therapy[1].

This paper makes two contributions. First, we present our survey results and ethnographic analysis. Second, we present a class of devices that use computer vision and decision theory to provide engaging arts activities for older adults and therapists. We present three instances of the device, and show some basic tests in a laboratory setting (without end users at this stage). 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 for handwashing assistance [Hoey *et al.*, 2008; Mihailidis *et al.*, 2008]. 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 [Mynatt *et al.*, 2000], the Assisted Cognition Project [Kautz *et al.*, 2002] and the Nursebot Project [Pineau *et al.*, 2003]. 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 [Pollack, 2006], one aspect of the Nursebot Project, applies a POMDP in the development of the planning and scheduling aspect of the system [Pineau *et al.*, 2003].

Partially observable Markov decision processes (POMDPs) provide a rich framework for planning under uncertainty[Åström, 1965]. 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 [Pineau *et al.*, 2003], spoken-dialog systems [Williams *et al.*, 2005] and vision-based systems for assistive technology [Hoey *et al.*, 2007], POMDPs can be used to optimize controllers that rely on the partial and noisy information provided by sensors such as sonars, lasers, video cameras and microphones.

A general POMDP model for assistive systems is presented in [Hoey *et al.*, 2005], in which the state space is broken into pieces relating to the *task*, the *user behaviour* and the *user internal states*. The model presented here specialises this model of general assistance for the task of facilitating a visual art activity. In this case, the *task* refers to specific interface situations, *behaviours* are a person's actions on the interface or visual gestures (e.g. facial expressions and gestures), and the *internal states* are the user's engagement and responsiveness.

Touchscreen devices for visual artworks have recently been developed [Raffle *et al.*, 2007], however not for the population of persons with dementia, and not as a tool for arts therapists specifically. We show one of our ePADs on a multi-touch display, touch screen devices that allow interaction at multiple points, by both hands or with objects such as brushes or pens. A variety of artistic activities can be produced on such displays [Ullmer and Ishii, 2000; Muise and Yim, 2008], including musical interfaces [Kaltenbrunner *et al.*, 2006].

3 User-Centered Design

Ethnography is the study of human practices and interactions with the environment [Macaulay *et al.*, 2000]. Ethnographers usually study their subjects through 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 [Macaulay *et al.*, 2000]. In this section, we review some key statistics from the survey results, followed by design ethnographic analysis of a survey of 133 practicing art therapists, and of input from a practicing art therapist who specialises in therapy for persons with dementia [Burns, 2009].

3.1 Survey Results

The respondents to the survey were from Canada (89), the UK (39) and from Ireland, Netherlands, Switzerland, Taiwan, Canada/USA and unknown (one each). 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.). There were 60 respondents who work with older adults, and 46 (77%) specifically work with people who have dementia. There were 28 respondents who said their survey was specifically about therapy for persons with dementia, geriatrics, older people in long-term care or Palliative care, or veterans. We will focus on this sub-group, and compare to the overall results. We also focus on the difference between UK and Canadian art therapists.

Of the activities that were specifically polled, arts therapists felt that older adults (both with and without dementia) most enjoyed painting (80%), drawing (63%), and sculpture (43%). Dance (30%), finger painting (27%), and writing (25%) were modestly ranked while music composition (18%) and theatre performances (8%) received the lowest scores. Other tasks, such as singing, collage, and constructing small objects (e.g. building a bird house) were often listed.

[1]However, note that in the case of persons with a progressive illness such as Alzheimer's, the artwork may be used more "as" therapy (where the art making process is the therapy) than "in" therapy (where the art is used as input for therapeutic analysis).

The survey asked questions relating to choice of activities and guidance. Therapists mostly felt that selecting the creative activity was a joint decision, 84% of the UK and 67% of the Canadian arts therapists stated their client selects their own activity. As shown in Figure 1, there is a significant difference between Canadian and UK therapists' guidance levels, with Canadian therapists giving more guidance than UK therapists. The figure also shows how persons with dementia require more guidance overall.

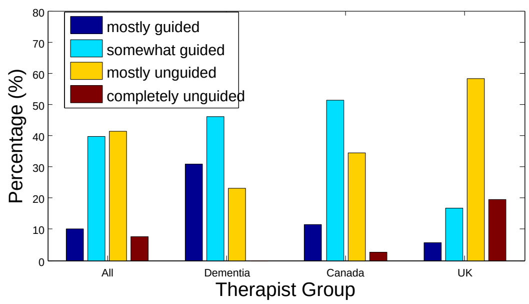

Figure 1: Guidance results from survey.

Overall, therapists felt most strongly that the device should be able to allow the client to select colours, be able to print and save copies of their art, and to use classic 'tools' (such as a paintbrush) when interacting with the device and were least in favour of the device auto-correcting shakey lines, playing background music, and performing auto-filling of a shape the user is trying to colour in. Therapists who worked with older adults with dementia had the same preferences.

Therapists from both nations had very similar rankings in their approach to engaging clients, with visual, verbal and gestural being the most commonly used techniques. Most (89%) therapists who worked with dementia stated that facial expression allowed them to tell if a person was engaged. All therapists felt that the device should be designed to be used personally as well as in a group.

Therapists commonly reported that they considered a successful outcome to be when their clients have a "smile on their face", demonstrate a "positive change", become more relaxed, or are generally engaged in the creative arts process. Lack of confidence, physical difficulties with materials and tools, and confusion and memory problems were commonly cited as difficulties or barriers older adults faced when attempting creative arts. Interesting, therapists who work with older adults with dementia felt more strongly than other groups that the device should be capable of having shared tasks with family members.

3.2 Design Ethnographic Analysis

The design ethnographic analysis found two major categories of implications, structural and design. Structural implications relate to the underlying abilities of the device in its functioning, whereas design implications relate to the perceived interface of the device. Four major structural (**S1–S4**) and three major design (**D1,D2,D3**) implications were uncovered, and are addressed by the device proposed in Section 4. Additional implications not yet addressed are reviewed in Section 6.

S1. Cutomisability
To be a useful assessment tool, the devices should be adaptable by the art therapist. This adaptability includes selection of activities and levels of difficulty. Beyond this, the device could even be configurable by the art therapist from a high level, allowing them to drop and drag activities and materials onto the project space according to the skills and preferences of the client as determined in previous evaluative sessions.

A further way in which art therapy is empowering to people is that it gives them an opportunity to learn new things. So while it is good to bear in mind the particular limitations of users with dementia, devices should handle the potential for skills acquisition. In other words, providing a challenge is not inherently a bad thing, and can in fact be empowering to users, even those with dementia.

S2. Adaptivity
The device should adapt to users. This is especially important in the case of art therapy for people with dementia, since the disease involves a progression through various stages of ability. In current practice, early evaluative sessions provide insight to art therapists about the appropriate level of choice for their clients, which is subsequently scaled accordingly. These sessions also provide insight into the individual preferences of clients (e.g. for materials). Bearing this in mind, the art therapy tool should be able to learn and to individualize materials, activities, and activity difficulty for clients.

The other benefit of developing an adaptive device is that it would provide useful feedback to art therapists. For example, when a client is having trouble with an activity, the device can automatically change the activity until a suitable level of difficulty is reached. It would also be helpful to then provide this graded assessment information to the therapists to show them their clients' progress (see implication **S4** below).

S3. Passivity
The dominant paradigm for art therapy in the UK is a non-directive approach whenever possible. Individuals with dementia may require a more directive (i.e. hands-on, therapist-led, guided) approach at times, though it is generally considered preferable to leave decision making in the hands of the clients. Such practice provides the art therapist with greater insight into the client's state of mind than if the therapist tells the client what to do. However, 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.

S4. Assessment
Computerised devices should provide quantitative measures to be used for assessment purposes. Current practice consists of refined qualitative measures for determining clients' individual abilities, needs, and goals. These qualitative measures are less effective when the art therapist is confronted with a new group of clients who each need an individual assessment. It would be a benefit to art therapists to have a tool that enabled them to quickly assess and evaluate the abilities and affective responses of new clients, in particular so that an individual's needs do not get lost in a group setting.

D1. Choice
A key element in empowering clients through art making is their ability to choose. Many people with dementia can feel disempowered in their day-to-day lives, having very few opportunities to make decisions and possibly feeling like no one

is listening to them. A unique benefit of art therapy, therefore, is that it promotes the person's ability to take control of his/her own art making, and to take ownership of what he/she has created. The art therapy tool, therefore, should aim to offer choice of activity, materials, and colors.

D2. Simplicity

In all art therapy cases, but arguably especially for clients with dementia, it is important to avoid overwhelming 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. It is recommended to start users with the most minimal possible starting point appropriate for their current level of ability.

D3. Touch

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.

4 ePAD Class

In this section, we present the ePAD class of devices for engaging users in art. Devices in this class combine sensing of a user with user modeling and decision making for optimisation of user engagement.

4.1 Device Overview

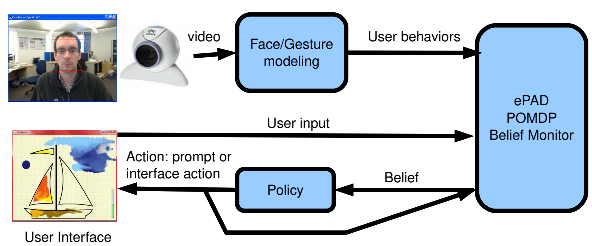

Figure 2: Overview of the system.

A diagram of an ePAD is given in figure 2. 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 camera provides video showing the user from some perspective, and a computer vision system recognises user behaviours (e.g. facial expressions or gestures). The recognised behaviors are used as input to a belief monitoring system that uses an instance of an ePAD POMDP model (see Section 4.2). This model maintains a belief about the user's current state, and passes this to a policy. The policy optimises utility over user engagement, and chooses an action for the system to take that it predicts will promote user engagement. This action is delivered to the interface, possibly through speakers if an audio prompt.

4.2 ePAD POMDP

An ePAD uses a probabilistic, decision theoretic model: a partially observable Markov decision process (POMDP) consisting 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 system actions cause stochastic state transitions, with different transitions being more or less rewarding (reflecting the relative utility of the states and the action costs). 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 (approximately) maximized.

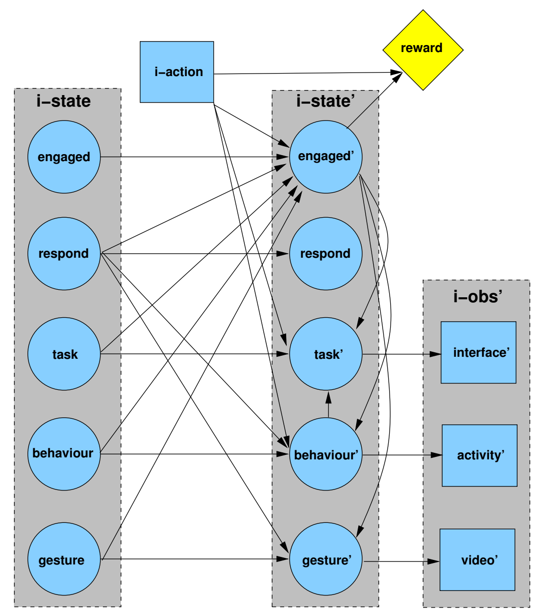

Figure 3: The ePAD POMDP model as a 2-time slice decision network: The full network can be obtained by unrolling in time. The primed variables are those that occur after an action. Observed variables are shown with a rectangular box, and are omitted for unprimed (pre-action) states for clarity.

ePAD defines a class of POMDP models, as shown as a Bayesian decision network in Figure 3. 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 (**i-action**) the system can take are to do nothing (a_0), give an audio prompt (a_p) or to modify the interface ($a_1, \ldots, a_M$, e.g. adding color, changing the activity) , where M is the number of possible modifications the system can make in the application. The ePAD class does not specify all the values of the variables and all the probabilistic relations, as described below. An ePAD instance is a POMDP for a particular art activity, allowing art therapists to customise the applications by adding new activities (implication **S1**, above), and only having to change certain aspects of the underlying model in well defined ways. We will see in Section 4.4 how multiple instances can be combined in a hierarchical model. The structure of the ePAD model, as in the COACH system [Mihailidis *et al.*, 2008], is a *passive* system: it only acts when necessary, allowing a user to do what they want, so long as they are engaged (implication **S3**).

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

1. The **behaviour**$\in\{$*interactive,active,inactive*$\}$ is whether the user is actively doing something on the interface, and is inferred from observations of their finger interactions (**activity** observation). The observation model $P(\mathbf{activity}|\mathbf{behaviour})$ is activity dependent. In a painting activity, interactive behaviours may be touching the screen in the region of a newly added colour.
2. The **gesture** variable is a set of gestures that indicate that a user is engaged with the device, and are inferred from the video stream. For example, the **gestures** could be gaze directions ($\{$*looking,not looking*$\}$), indicating if a person is looking at the screen. The **gestures** and the observation model $P(\mathbf{video}|\mathbf{gesture})$ are activity dependent.
3. The $\mathbf{task} \in \{t_1, t_2, \ldots, t_N\}$ denotes the stage of the current application at which the user is operating. For example, in a collage application, the states t_i could indicate how many items are left to stick on the collage. This variable is less important in an art application than it is in, for example, the handwashing assistant [Hoey *et al.*, 2007]. However, it is still necessary so the system can infer when an activity is completed (e.g. there is nothing left for a user to do).

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

1. The user's **engagement**$\in\{$yes,confused,no$\}$ is the key element of this model, as maintaining engagement is the primary purpose of the device. A user can be engaged (*yes*), or disengaged (*confused*,*no*). The dynamics of engagement, and the effects of the engagement on a user's behaviours, gestures, and task are activity dependent, but are modifiable with a small number of parameters, as described below.
2. We also model the user's responsiveness to the system's actions, **respond**. This variable is actually factored into a number of variables for each task: $\{\mathbf{respond\_cue}, \mathbf{respond\_a_1}, \ldots, \mathbf{respond\_a_M}\}$, giving a user's responsiveness to an audio cue and to each of the interface actions $a_1, \ldots, a_M$, respectively. These variables are also activity dependent, but again, we can define general classes of responsiveness (e.g. to audio prompts).

These user modeling variables are designed to address the structural design implication of adaptivity in **S2** above, by defining a range of user types (e.g. responsive/engaged). For each user type, the model dynamics will give rise to a different strategy on the part of the device, allowing for adaptivity as the user interacts. Use of this model by monitoring the devices *belief* in the user's state allows for a direct and quantitative measure of a user's ability, engagement, and responsiveness in a particular application (implication **S4**).

The dynamics of the POMDP hinges on the user's **engagement**. For example, likelihood of interaction with the device if responsive is given by $P(\mathbf{behaviour'}|\mathbf{behaviour}, \mathbf{engaged'}, \mathbf{respond})$ 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 at the screen, and the system does something of interest, 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 based solely on the user's engagement, with +10,-1,-2 if the user is engaged, confused or not engaged, respectively. Motivational prompts are slightly costly (cost of 0.5), but only if the user is engaged. This models the effect of a prompt reducing feelings of independence in a user if they are already engaged (**S3**). Note that this cost is separate from the indirect costs incurred if a user is prompted when engaged, possibly leading to confusion.

4.3 Policy Computation

To compute an approximate policy, we used the SymbolicPerseus package [Poupart, 2005][2]. It implements a point-based approximate solution technique based on the Perseus algorithm [Spaan and Vlassis, 2005] combined with Algebraic Decision Diagrams as the underlying data structure.

4.4 Hierarchical Extension

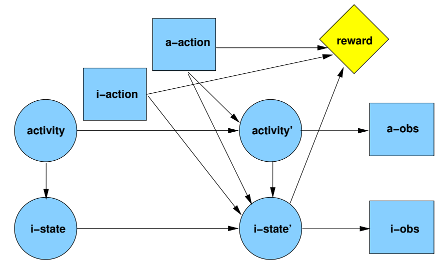

Figure 4: The hi-ePAD hi-level POMDP model as a 2-time slice decision network. The **a-action** selects among a number of activities. Observations are not shown for pre-action (unprimed) states for clarity.

The ePAD model described above is for a particular activity. However, we would ideally like the device to be able to switch between activities, on the initiative of user, therapist, or system. This ability will allow for seamless integration of new applications (e.g. by art therapists) on a working device (implication **S1**). To do this, we group all factors in the ePAD model (Figure 3) into a single factor **i-state**, and all observations into **i-obs**, and call the actions the system can take for each activity **i-actions**. We then add a new variable **activity**, a new observation **a-obs**, and a new set of system actions **a-action**, as shown in Figure 4. The **activity** variable defines the activity currently being performed, the **a-action** is to do nothing (allowing the currently running activity to proceed),

[2] code available at www.cs.uwaterloo.ca/~ppoupart/software

or to change to a new activity. The **a-obs** is for user initiated switches of activities.

This type of hierarchical POMDP poses no conceptual difficulties for the ePADs, but does increase the size of the state space significantly. Currently, we use a "flattened" version of this model (so planning occurs over the entire state space), but hierarchical planning algorithms will be applied in the future [Theocharous *et al.*, 2004]. Note that the model in Figure 4 is a constrained version of a full hierarchical POMDP, as the activities cannot terminate without an explicit action (**a-action**) or observation (**a-obs**).

5 Example ePAD Instances

This section presents three example ePAD instantiations. The first one goes into some detail and gives an example interaction. The second and third give brief overviews.

5.1 ePAD I: Finger painting on vertical touchscreen

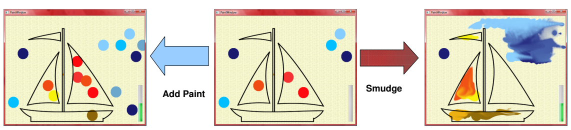

Figure 5: The application and the actions which can be performed upon it. The only user action is to mix the colours. The system can also change the background line drawing.

Our first ePAD involves a painting application in which blobs of color are placed on a canvas with a simple background (line drawing). The user can mix the colors with a finger, with realistic blending effects. 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 screen or by changing the background image. The hardware is a NEC Multisync LCD 2010X touchscreen running on a standard desktop, with a Logitec Quick Cam Pro 9000 mounted above the screen. Figure 5 shows examples of the interface.

The device monitors a person's face, to estimate their engagement (design implication **D3**). The OpenCV [Bradski and Kaehler, 2008] implementation of the Viola-Jones face detector [Viola and Jones, 2001] is used to detect the presence of a frontal face. We present a simple demonstration of the system in the laboratory with a single non-demented subject acting according to the situations described. These are meant only to illustrate the functionality of the device and user monitoring, not as end-user tests.

The example in Figure 6 shows a person who is responsive to system actions and who is engaged in the task. 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 (N). The system adds a blob of colour, and observes the the user "plays" by touching the added blob of colour. The user is estimated to be interactive (I), and more responsive to colours, and more engaged.

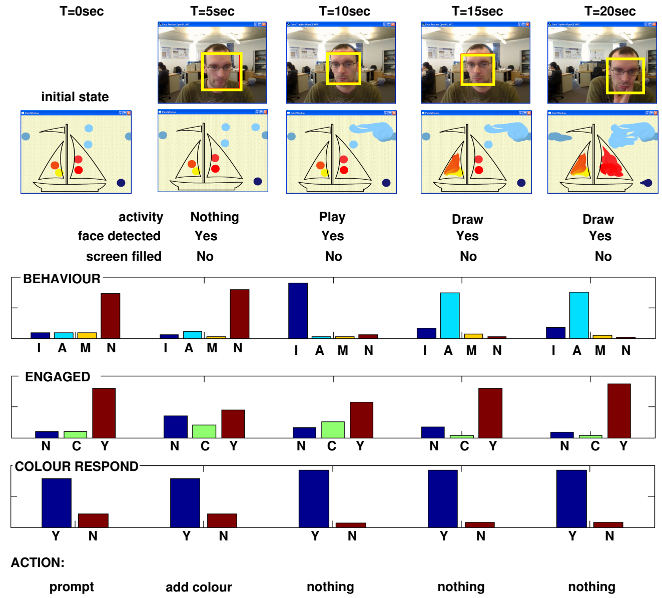

Figure 6: Example sequence of an engaged user. Behaviour can be interactive (I), active (A), intermittent (M) and inactive (N). Engaged can be no (N), confused (C) and yes (Y). Colour Respond can be yes (Y) or no (N). The observations are shown below the images of the interface. The actions taken by the system are along the bottom. Time runs from left to right, in 5 second increments. The marginal beliefs (in $[0, 1]$) are shown for each variable as bar plots.

The person then continues using the system increasing both the belief in engagement = Y and behaviour = active. When the person is engaged the system takes no action.

5.2 ePAD II: Collage on Multi-Touch Surface

Figure 7: Collage on a Multi-touch table

A second application involves collage of multi-colored objects (leaves) onto a background picture (a tree), shown in Figure 7. The user can drag the objects around and "glue" them down using a double tap. There are two buttons that allow a user to drop random collage pieces on the surface ("drop"), and to get rid of un-stuck items ("brush"). The user behaviors are drag, glue, drop, or brush. Dragging new objects is considered "interactive" behaviour as in the painting application. The task models how much of the screen is filled, as in the finger painting application. The actions of the system can be to drop or brush, or give a motivational prompt. The user's responsiveness to all three actions are modeled. Again, a simple web camera detects if a person is looking at the screen. In this case, however, the camera positioning is more difficult, and additional uncertainty is added to the observation function $P(video|gesture)$.

The display is a FTIR (frustrated total internal reflection) multi-touch surface. A large perspex screen is physically embedded into a custom-built coffee table. A computer inside the table runs all the software. Graphics are rear projected using a internal short throw projector onto the screen. Infra red lights shine from the side of the screen, which is covered by a thin layer of silicone. When a finger or another object makes contact with the surface, the light is directed into the table to an infra red camera. This ePAD can therefore detect multiple fingers and objects when they come into contact with the screen, even allowing multiple people (e.g. a therapist and client together) to use the ePAD simultaneously.

5.3 ePAD III: Flipbook on Blade Tablet

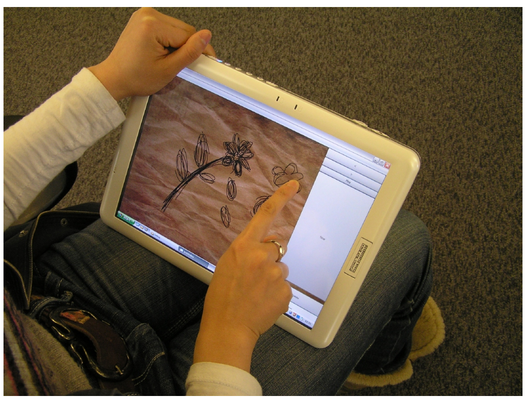

Figure 8: Drawing on a Blade Tablet PC.

Our third application is a flipbook animation, shown in Figure 8. A user can create and review simple animations by freehand line drawing. Buttons are for adding a new frame, changing from a pen to an eraser, playing and stopping the animation. The user behaviors are to draw/erase, to add a frame, or to play/stop. This application is our latest prototype, and is still in development. Possible system's actions are to suggest drawings, to highlight buttons, or to give prompts, but the dynamic nature of the artwork makes the precise modeling of the effects of system actions more involved. The task may involve number of frames completed, or completion of each individual frame. These task states require some additional complexity due to the dynamic nature of the art work.

The display for this ePAD is a PaceBlade^{TM} touch tablet. This small device is very portable, and has a sensitive single-touch screen that can be used with a finger or a pen. The advantages of this tablet are its portability and sensitivity, and the main disadvantage is size. Again, we use a Logitech webcam with face tracking. Additional uncertainty in the observation function takes into account the mobility of the device.

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 art 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 somehow incorporate the tactile into the activity, and to link the art materials to the sounds they would naturally make, e.g. linking 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's 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, including level of difficulty. Our work will address this by allowing arts therapists to program the device it 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 backend. 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 targetted focus groups, followed by user testing with wizard-of-oz studies.

D7. Feedback
Knowing what activities the client did with the device was ranked as the most important feedback for the therapists, followed by the number of times they used the device, who else was using the device, and the time of day the device was used. 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 touchscreen activity that the art therapist can play back.

Feedback to art therapists should ideally be presented in a visual manner, as art therapists are highly visual people. And as they deal almost exclusively with qualitative data in their practice, it is important to avoid outputting data in the form of complicated numbers and graphs. A proposed solution is to present feedback in the form of several sliding scales, which would offer a visual of the client's position on a continuum. For example, one such continuum might be degree of difficulty of activity; another might be level of engagement; another might be degree of frustration.

The other useful feedback for art therapists would be a representation of the pattern of the client. Specifically, it would help to show whether the client has improved in certain areas, attained new skills, or declined in ability. 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.

Acknowledgements
The authors thank Laurel Young and all survey participants, as well as Annalu Waller and Rolf Black for help with the Blade. This work was supported by the American Alzheimer's Association under the *Everyday Technologies for Alzheimer's Care* program.

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 a progressive illness such as Alzheimer's disease. 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, 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.

References

[Åström, 1965] K. J. Åström. Optimal control of Markov decision processes with incomplete state estimation. *Journal of Mathematical Analysis and Applications*, 10:174–205, 1965.

[Bradski and Kaehler, 2008] Gary Bradski and Adrian Kaehler. *Learning OpenCV*. O'Reilly Media, 2008.

[Burns, 2009] Jane Burns. *An interpretive Description of the patterns of practice of art therapists working with older people who have dementia in the UK*. PhD thesis, Queen Margaret University, Edinburgh, Scotland, 2009.

[Geda *et al.*, 2009] Yonas E. Geda, Rosebud Roberts, David Knopman, Vernon S. Pankratz, Teresa Christianson, Bradley Boeve, Robert Ivnik, Eric G. Tangalos, Walter Rocca, and Ronald Petersen. 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*, 2009.

[Harlan, 1990] 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.

[Hoey *et al.*, 2005] Jesse Hoey, Pascal Poupart, Craig Boutilier, and Alex Mihailidis. POMDP models for assistive technology. In *Proc. AAAI Fall Symposium on Caring Machines: AI in Eldercare*, 2005.

[Hoey *et al.*, 2007] Jesse Hoey, Axel von Bertoldi, Pascal Poupart, and Alex Mihailidis. Assisting persons with dementia during handwashing using a partially observable Markov decision process. In *Proceedings of the International Conference on Vision Systems*, March, 2007. http://biecoll.ub.uni-bielefeld.de/volltexte/2007/12.

[Hoey *et al.*, 2008] Jesse Hoey, Pascal Poupart, Axel von Bertoldi, Tammy Craig, Craig Boutilier, and Alex Mihailidis. Automated handwashing assistance for persons with dementia using video and a partially observable markov decision process. *To appear: Computer Vision and Image Understanding available from www.computing.dundee.ac.uk/staff/jessehoey/blearn/*, 2008.

[Kaltenbrunner *et al.*, 2006] Martin Kaltenbrunner, Sergi Jorda, Gunter Geiger, and Marcos Alons. The reactable*: A collaborative musical instrument. In *15th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises*, pages 406–411, 2006.

[Karp *et al.*, 2006] Anita Karp, Stephanie Paillard-Borg, Hui-Xin Wang, Merril Silverstein, Bengt Winblad, and Laura 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.

[Kautz *et al.*, 2002] 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.

[Macaulay *et al.*, 2000] Catriona Macaulay, D. Benyon, and C. Crerar. Ethnography, theory and design: From intuition to insight. *International Journal of Human-Computer Studies*, 53:35–60, 2000.

[Mihailidis *et al.*, 2008] Alex Mihailidis, Jen Boger, Marcelle Candido, and Jesse Hoey. The coach prompting system to assist older adults with dementia through handwashing: An efficacy study. *BMC Geriatrics*, 8(28), 2008.

[Muise and Yim, 2008] K. Muise and J. Yim. A touching harmony: Midas in artistic practice. In ACM, editor, *Proceedings of the 2nd international Conference on Tangible and Embedded interaction*, pages 213–214, Bonn, Germany, Feb 2008. ACM.

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

[Pineau *et al.*, 2003] 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.

[Pollack, 2006] Martha 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.

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

[Raffle *et al.*, 2007] Hayes Raffle, Cati Vaucelle, Ruibing Wang, and Hiroshi Ishii. Jabberstamp: embedding sound and voice in traditional drawings. In *Proceedings of the Conference on Interaction Design and Children (IDC)*, Aalborg, Denmark, June 2007. ACM Press.

[Rusted *et al.*, 2006] Jennifer Rusted, Linda Sheppard, and Diane 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.

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

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

[Theocharous *et al.*, 2004] Georgios Theocharous, Kevin Murphy, and Leslie Kaelbling. Representing hierarchical POMDPs as DBNs for multi-scale robot localization. In *Proc. International Conference on Robotics and Automation (ICRA 2004)*, 2004.

[Ullmer and Ishii, 2000] B. Ullmer and H. Ishii. Emerging frameworks for tangible user interfaces. *IBM Systems Journal*, 39(3-4):915–931, July 2000.

[Viola and Jones, 2001] 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.

[Williams *et al.*, 2005] 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, August 2005.

[Wilson *et al.*, 2002] Robert S. Wilson, Carlos F. Mendes de Leon, Lisa L. Barnes, Julie A. Schneider, Julia L. Bienias, Denis A. Evans, and David A. Bennett. Participation in cognitively stimulating activities and risk of incident Alzheimer disease. *Journal of the American Medical Association*, 287(6):742–748, February 2002.