

Affectively Aligned Cognitive Assistance Using Bayesian Affect Control Theory

Luyuan Lin¹, Stephen Czarnuch², Aarti Malhotra¹, Lifei Yu¹,
Tobias Schröder³, and Jesse Hoey¹

¹ School of Computer Science, University of Waterloo, Waterloo, Ontario, N2L3G1

² Institute of Biomaterials and Biomedical Engineering,

University of Toronto, 500 University Ave., Toronto, Ontario M5G1V7

³ Potsdam University of Applied Sciences, 14469 Potsdam, Germany

Abstract. This paper describes a novel emotionally intelligent cognitive assistant to engage and help older adults with Alzheimer’s disease (AD) to complete activities of daily living (ADL) more independently. Our new system combines two research streams. First, the development of cognitive assistants with artificially intelligent controllers using partially observable Markov decision processes (POMDPs). Second, a model of the dynamics of emotion and identity called *Affect Control Theory* that arises from the sociological literature on culturally shared sentiments. We present background material on both of these research streams, and then demonstrate a prototype assistive technology that combines the two. We discuss the affective reasoning, the probabilistic and decision-theoretic reasoning, the computer-vision based activity monitoring, the embodied prompting, and we show results in proof-of-concept tests.

1 Introduction

Persons with dementia (PwD, e.g. Alzheimer’s disease) have difficulty completing activities of daily living, such as handwashing, preparing food and dressing. The short-term memory impairment that is a hallmark of Alzheimer’s disease leaves sufferers unable to recall what step to do next, or what important objects look like, for example. We have been developing a smart home system called the *COACH* to assist older adults with dementia to carry out basic ADL (e.g. handwashing) through step-by-step audiovisual prompts [2,7,14]. The *COACH* is effective at monitoring and making decisions about when/what to prompt [14], and works well for some persons, but not as well for others. Considering the heterogeneity in socio-cultural and personal affective identities, a primary reason for lack of effectiveness may be the static, non-adaptive nature of the “canned” (pre-recorded) prompts. While we have made significant effort to design prompts founded on the methods and styles of human caregivers [19], a simple “one size fits all” style of prompting may be limiting. For example, one person might find our prompts to be too imperious, and would respond better to a more servile approach. However, this will not be the case with every person, and some may prefer the more imperative prompting style. Each person comes from a different

background, has a different sense of “self”, and has different emotional responses to prompts, whether given by human or machine. Affective identity is believed to be a powerful tool for reasoning about illness in general [11]. In particular, studies of identity in Alzheimer’s disease have found that identity changes dramatically over the course of the disease [15], and that persons with AD have more vague or abstract notions of their identity [17].

In this paper, we build explicit models of emotional identity and personality into a cognitive assistant, and we give ideas about how this can improve the overall (cross-individual) effectiveness and potential uptake of such systems. Our system use a RGB-D camera to detect the body postures of a person while handwashing, and infers both the functional meaning (e.g. does the person have soap on their hands or not, or, is the water running?), and the affective meaning (e.g. is the person feeling powerless, in control, angry, or depressed?) of the observed behaviours. The observed functional and affective behaviours are then fed into a reasoning engine that uses a partially observable Markov decision process (POMDP), a probabilistic and decision theoretic model of both the handwashing task and the affective identity of the person using the system. The affective component is based upon a sociological theory called *Affect Control Theory* [5], which models the dynamic affective identities and behaviours of the person and the handwashing assistant. The POMDP policy produces an approximately optimal action for the system to take, again on both functional (e.g. what step is next) and affective (e.g. imperious vs. servile delivery of the prompt) dimensions. This prompt is delivered as a video of an embodied caregiver acting in a style that is consistent with the recommended affective action.

2 Background

2.1 Affect Control Theory

Affect Control Theory (ACT) arises from work on the sociology of human interaction [5]. ACT proposes that social perceptions, behaviours, and emotions are guided by a psychological need to minimize the differences between culturally shared fundamental affective sentiments about social situations and the transient impressions resulting from the interactions between elements within those situations. Fundamental sentiments, \mathbf{f} , are representations of social objects, such as interactants’ identities and behaviours or environmental settings, as vectors in a three-dimensional affective space. The basis vectors of the affective space are called Evaluation/valence, Potency/control, and Activity/arousal (EPA). The EPA space is hypothesized to be a universal organizing principle of human socio-emotional experience, based on the discovery that these dimensions structure the semantic relations of linguistic concepts across languages and cultures [16]. They also emerged from statistical analyses of the co-occurrence of a large variety of physiological, facial, gestural, and cognitive features of emotional experience [4], relate to the universal dimensionality of personality, non-verbal behaviour, and

social cognition [18], and are believed to correspond to the fundamental logic of social exchange and group coordination [18].

EPA profiles of concepts can be measured with the *semantic differential*, a survey technique where respondents rate affective meanings of concepts on numerical scales. In general, within-cultural agreement about EPA meanings of social concepts is high even across subgroups of society, and cultural-average EPA ratings from as little as a few dozen survey participants are extremely stable over extended periods of time [6]. For example, the EPA for the identity of “nurse” is [1.65, 0.93, 0.34], meaning that nurses are seen as quite good (E), a bit powerful (P), and a bit active (A)¹. Comparatively a “patient” is seen as [0.9, −0.69, −1.05], less powerful and less active than a “nurse”. Social events cause transient impressions, τ , of identities and behaviours that deviate from their corresponding fundamental sentiments, \mathbf{f} . ACT models this formation of impressions from events with a minimalist grammar of the form agent-behaviour-client. Consider, for example, a nurse (agent) who ignores (behaviour) a patient (client). Observers agree, and ACT predicts, that this nurse appears (τ) less nice (E), and less potent (P), than the cultural average (\mathbf{f}) of a nurse. The Euclidean distance between τ and \mathbf{f} is called the *deflection* (D), and is hypothesized to correspond to an aversive state of mind that humans seek to avoid (the *affect control principle*). For example, the nurse who “ignores” a patient has a deflection of over 15 (very high), whereas if the nurse “comforts” the patient, the deflection is 1.5 (very low). The *affect control principle* also allows ACT to compute *normative* actions for artificial agents: those that minimize deflection. ACT has been shown to be a powerful predictor of human behaviour [12].

2.2 Partially Observable Markov Decision Processes

A partially observable Markov decision process (POMDP) [1] is a general model of stochastic control that has been extensively studied in operations research and in artificial intelligence. A POMDP consists of a finite set \mathcal{S} of states; a finite set \mathcal{A} of actions; a stochastic transition model $\Pr : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$, with $\Pr(s'|s, a)$ denoting the probability of moving from state s to s' when action a is taken², and $\Delta(\mathcal{S})$ is a distribution over \mathcal{S} ; a finite observation set Ω ; a stochastic observation model with $\Pr(\omega|s)$ denoting the probability of making observation ω while the system is in state s ; and a reward assigning $R(a, s')$ to a transition to s' induced by action a . A *policy* maps *belief states* (i.e., distributions over \mathcal{S}) into choices of actions, such that the expected discounted sum of rewards is (approximately) maximised. In this paper, we will be dealing with *factored* POMDPs in which the state is represented by the cross-product of a set of variables or features. Assignment of a value to each variable thus constitutes a state. POMDPs have been used as models for many human-interactive domains, including intelligent tutoring systems, and human assistance systems [7].

¹ EPA values range from −4.3 to 4.3 by convention.

² primes indicate post-event variables, unprimed are pre-event variables.

2.3 BayesAct

Recently, we have developed a probabilistic and decision-theoretic generalization of ACT, and have demonstrated how it can be leveraged to build affectively intelligent artificial agents [9]. The new model, called *BayesAct*, can maintain multiple hypotheses about sentiments simultaneously as a probability distribution, and can make use of an explicit utility function to make value-directed action choices. This allows the model to generate affectively intelligent interactions with people by learning about their identity, predicting their behaviours using the affect control principle, and taking actions that are simultaneously goal-directed and affect-sensitive.

A *BayesAct* POMDP models an interaction between two agents (human or machine) denoted *agent* and *client*. The state is the product of six 3-dimensional continuous random variables corresponding to fundamental and transient sentiments about the *agent's* identity ($\mathbf{F}_a, \mathbf{T}_a$), behaviour ($\mathbf{F}_b, \mathbf{T}_b$) and *client's* identity ($\mathbf{F}_c, \mathbf{T}_c$). The transient impressions, $\mathbf{T} = \{\mathbf{T}_a, \mathbf{T}_b, \mathbf{T}_c\}$, evolve according to the deterministic impression-formation operator in ACT. Fundamental sentiments, $\mathbf{F} = \{\mathbf{F}_a, \mathbf{F}_b, \mathbf{F}_c\}$, are expected to stay approximately constant over time, but are subject to random drift (with noise Σ_f) and are expected to also remain close to the transient impressions because of the *affect control principle*. This allows us to estimate the posterior probability distribution over sentiments, $Pr(\mathbf{f}', \boldsymbol{\tau}' | \mathbf{f}, \boldsymbol{\tau})$, which gives the normative (expected) action as \mathbf{f}'_b . Thus, by integrating over $\mathbf{f}'_a, \mathbf{f}'_c, \boldsymbol{\tau}'$ and the previous state, we obtain a probability distribution over \mathbf{f}'_b that acts as a *normative action prediction*: it tells the agent what to expect from other agents, and what action is expected from it. This normative action is used as the POMDP policy directly.

BayesAct includes an application-specific set of random variables \mathbf{X} that are interpreted as the remainder of the state space, including *non-affective* elements of the domain (e.g. steps of the handwashing task). The dynamics of \mathbf{X} are application specific, but depend in general on the *deflection*, and on the propositional component of the action, a (which complements the affective component, \mathbf{b}_a , a 3D EPA vector). Finally, *BayesAct* has a two observation variables, Ω_x and Ω_f , that give evidence for the variables, \mathbf{X} and \mathbf{F}_b , respectively, through observation functions $Pr(\Omega_x | \mathbf{X})$ and $Pr(\Omega_f | \mathbf{F}_b)$.

3 System Description

3.1 Handwashing POMDP with Affective Reasoning

We use a model of the handwashing system with 8 *plansteps* corresponding to the different steps of handwashing, describing the state of the tap (on/off), and hands (dirty/soapy/clean and wet/dry). An eight-valued variable *PS* describes the current planstep. The client's behaviour is modeled with a six-valued variable *BEH* describing his/her actions: turn on/off water, use soap, use towel, rinse and null (do nothing). There are probabilistic transitions between plansteps described in a probabilistic plan-graph (e.g. a PwD sometimes uses soap first, but

sometimes turns on the tap first). We also use a binary variable AW describing if the PwD is *aware* or not. In [14], we also had a variable describing how *responsive* a person is to a prompt. Here, we replace that with the current deflection in the interaction. Thus, $\mathbf{X} = \{PS, BEH, AW\}$ and the dynamics of the PS are

- If the client is aware, then if there is no prompt from the agent, the client will advance stochastically to the next planstep with a probability that is dependent on the current observation of client behaviour and the deflection, D . If the client does not advance, she loses awareness.
- If the client is aware and is prompted, and D is high, then the prompt will likely confuse the client and (stochastically) cause him/her to lose awareness.
- If the client is not aware, then if there is a prompt from the agent, and D is low, the client will likely follow the prompt and gain awareness. Otherwise (i.e. there is no prompt, or the deflection D is high), the client will not do anything (or do something other than the one prompted) with high probability.

We have found that a fixed affective policy may work well for some affective identities, but not for others, whereas the actions suggested by *BayesAct* work well across the different identities that the client may have [8].

3.2 Functional Motion Classification

Caregivers of older adults with dementia have indicated that any assistive technologies must integrate into the environment to reduce the likelihood of stigmatization, but be out of reach [2]. The tracker we use classifies individual body parts from a single overhead depth image on a per-frame basis [3]. The tracker first uses a random decision forest with a simple depth feature to provide intermediate multiclass probability density functions (PDF) for each sampled image pixel. The tracker then proposes final body part positions by aggregating the information contained in the underlying PDF. The tracker is trained on a set of images that are manually annotated to optimize key parameters. The optimal parameters are then used to train a final decision forest resulting in a new depth-based hand tracker. The tracker outputs the locations of the two hands and the head, and has been independently evaluated [3].

The hand locations are mapped to a set of pre-defined spatial regions (soap, tap, sink, water, towel). If multiple areas are detected, then a set of rules, based on the distances from the region centers and current hand-locations, are applied to decide the “winner” region that is used as the observation $\omega_{\mathbf{x}}$ (of variable \mathbf{x}) in the POMDP. Further details are found in [10].

3.3 Affective Motion Classification

Hands’ coordinates obtained from the hand tracker are used to extract EPA values as follows. Evaluation (“E”) stays neutral for all situations as it is the most difficult one to measure (e.g. facial expressions or vocal tone could be used in future). The mean of the distances between the user’s two hands within a set

of $n = 10$ frames is an indication of the “open-ness” of a body posture, known to be a good indication of feelings of dominance or power. A piecewise linear function maps from this average inter-hand distance to potency, $P \in [-4.3, 4.3]$.

The activity (A) is based on the speed of movement of the user’s hands. In each pair of successive frames, the maximum difference between any two hand positions is computed. These differences are averaged over $n = 10$ frames, and a second linear interpolation function is used to map these differences to activity, $A \in [-4.3, 4.3]$. The linear interpolants for P and A are learned from experimental data. Further details and precise settings can be found in [10].

The EPA vector that results from the calculations above is used as the observation ω_f (of \mathbf{f}_b) in *BayesAct*. We set the covariance in the observation function $Pr(\omega_f|\mathbf{f}_b)$ to be such that the “E” dimension is ignored (infinite variance), and the “P” and “A” dimensions have relatively small effects on the agent’s estimate of \mathbf{f}_b . That is, we set the covariance to be a diagonal with entries (10000, 1.0, 0.5). The variance in “P” is set to be larger than that in “A” since the distance between user’s hands is a much weaker indication of potency than the speed of user’s hands is an indication of activity. If other measures were used for EPA observations, these variances could be adjusted accordingly.

3.4 Affective Prompting

We created a set of audio-visual prompts using a virtual human developed with the USC Virtual Human Toolkit (VHT)³. We built a set of six audio-visual prompts with five different emotional deliveries (e.g. “bossy”, “motherly” or “bored” - see screenshots in Table 1). An online survey was then conducted in which participants were asked to watch the 30 videos and rate them based on Evaluation, Potency, and Activity dimensions (on a discrete scale of -4 to +4 with increments of 1 for a total of 9 options). Following [6], we showed sets of concepts at either end of the scales: bad/awful to good/nice (Evaluation); impotent/powerless/little to potent/powerful/big (Potency); inactive/slow/quiet to active/fast/noisy (Activity). The questions were presented in randomized order. There were total of 27 respondents.

To determine consensus amongst participants, we followed the culture-as-consensus model measuring the shared knowledge of the culture within the respondents [6]. The method computes the Eigenvalues of the covariance matrix of all responses for each of E,P,A separately. These eigenvalues indicate the extent to which respondents agree in their ratings across all items. If the ratio of the first to second Eigenvalue is large, this reflects cultural commonality in the respondent’s ratings and provides evidence of one dominant factor governing respondent’s judgement [6]. The Eigenvalue ratios for E were 8.518, that for P was 1.523 and that for A was 1.914, indicating that the respondents agreed most on the evaluation dimension, with reasonable agreement on potency and evaluation (see details in [13]).

³ <https://vhtoolkit.ict.usc.edu>

For use in the handwashing assistant, we compute the mean value across all respondents (ignoring missing values). Given the propositional and emotional descriptions of desired prompts (obtained from the *BayesAct* POMDP), our system selects the video prompt that matches the propositional label and whose mean value (from the survey results) is closest in the EPA space.

4 Experiments and Results

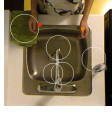

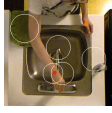


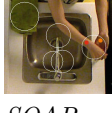
The system operates with server-stubs and client-stubs, and Google’s protocol buffer mechanism were used as the way to define the request and response messages shared by the two communicating parties. Open source libraries, such as zeromq and libVLC for prompts, were utilized as well. The experiments were conducted on a PC running 64-bit Ubuntu 12.04 LTS, with AMD FX(tm)-6300 Six-Core Processor 6 and Gallium 0.4 on llvmpipe (LLVM 0x300) Graphics. A kinect camera was mounted above the sink area and was the only sensor. Further details on the experiments can be found in [10].

Table 1 shows the first four steps of an example run with the system, along with the tracked hand locations, the prompts given and the various beliefs the system has about the user. We can see that the behaviours of the user are monitored correctly, and the system responds with appropriate prompts.

We did two tests with an actor washing her hands while the system observes and assists her in real time. The actor acted more powerfully (with her hands more “open”) and more actively (with her hands moving more quickly) in the first test than in the second one. \mathbf{f}_c^0 was set to $[1.61, 0.84, -0.87]$ in test #1, and was set to $[-0.64, -0.43, -1.81]$ in test #2⁴. Recall that \mathbf{f}_c denotes the agent’s belief of the client’s identity, and \mathbf{f}_c^0 denotes the initial value of this belief. Throughout the tests, the user behaviours in the first test generally had larger P and larger A values than those in the second test. The P and A values computed for user behaviours in test #1 reached an average of $[1.32, -1.3]$, while that in the second test was $[0.77, -1.74]$. Further, the f_c ’s in the first test generally had larger P and larger A values than those in the second test, and the system prompts in the first test generally had smaller P and higher A values. The mean of the EPA values of f_c ’s in the two tests were $[2.8, 1.03, -0.73]$ and $[1.13, -0.43, -1.47]$, respectively. And the mean of the EPA values of system prompts in the two tests were $[1.62, 0.32, 0.75]$ and $[1.53, 0.66, 0.08]$. Prompts with lower P values and higher A values are produced for identities with higher P values and higher A values. This correlation makes sense, since people who think of themselves as powerful persons tend to expect respect from others in interactions (i.e. prompts should be expressed to them with low potency levels), and that active people are likely to interact better with persons who are active as well — these “intuitions” are born out with *BayesAct* simulations, and thus are in accord with the predictions of Affect Control Theory.

⁴ These EPAs are close to the identities of “elder” and “lonesome elder”, resp.

Table 1. First 4 steps of an experimental test run with acted behaviours. \mathbf{f}_c : system’s estimate of the user’s affective identity. \mathbf{f}_b : user’s affective behaviour. \mathbf{b}_a : affective delivery of prompt. Plansteps are tap/hands/wet = [off/dirty/dry, on/dirty/dry, off/soapy/dry, on/soapy/dry, on/clean/wet, off/clean/wet, on/clean/dry, off/clean/dry]. Propositional prompts (prop) are 0:none, 1:water on, 2: soap.

time (sec)	\mathbf{f}_c (E,P,A)	behaviour		planstep belief/value	prompt		
		video/value	\mathbf{f}_b		prop	\mathbf{b}_a	avatar
0	$\begin{bmatrix} 1.70 \\ 1.40 \\ -1.39 \end{bmatrix}$	 <i>TOWEL</i>	$\begin{bmatrix} 0 \\ 1.86 \\ -1.7 \end{bmatrix}$	$\begin{bmatrix} 1 & 0 & 0 & 0 \dots \\ 0 & 0 & 0 & 0 \end{bmatrix}$ most likely: 0	1	$\begin{bmatrix} 1.82 \\ 0.22 \\ 0.47 \end{bmatrix}$	 “Hello I am so glad to have you here. Please turn on the water”
4	$\begin{bmatrix} 2.73 \\ 1.14 \\ -1.03 \end{bmatrix}$	 <i>TAP</i>	$\begin{bmatrix} 0 \\ 1.68 \\ -0.58 \end{bmatrix}$	$\begin{bmatrix} .26 & .74 & 0 & 0 \dots \\ 0 & 0 & 0 & 0 \end{bmatrix}$ most likely: 1	0	$\begin{bmatrix} - \\ - \\ - \end{bmatrix}$	
6	$\begin{bmatrix} 2.67 \\ 1.21 \\ -0.72 \end{bmatrix}$	 <i>RINSE</i>	$\begin{bmatrix} 0 \\ 1.49 \\ -0.16 \end{bmatrix}$	$\begin{bmatrix} .27 & .73 & 0 & 0 \dots \\ 0 & 0 & 0 & 0 \end{bmatrix}$ most likely: 1	2	$\begin{bmatrix} 1.51 \\ 0.12 \\ 0.52 \end{bmatrix}$	 “You are washing your hands. Please use the soap.”
10	$\begin{bmatrix} 2.57 \\ 0.69 \\ -0.66 \end{bmatrix}$	 <i>SOAP</i>	$\begin{bmatrix} 0 \\ 0.73 \\ -1.52 \end{bmatrix}$	$\begin{bmatrix} 0 & .01 & .35 & .64 \dots \\ 0 & 0 & 0 & 0 \end{bmatrix}$ most likely: 3	0	$\begin{bmatrix} - \\ - \\ - \end{bmatrix}$	

5 Conclusion

We have presented a prototype of an assistance system that reasons about affective identities. Our hypothesis is that older adults with AD will be more engaged with, and will adhere to more prompts by, the prompting system that uses affective reasoning. This will result in an increase in the number of ADL steps completed independently. In future, we plan to investigate notions of identity in Alzheimer’s disease, to work on measurement of EPA from verbal and non-verbal behaviours, and to develop dynamic (continuous) video prompts.

References

1. Åström, K.J.: Optimal control of Markov decision processes with incomplete state estimation. *J. Math. Anal. App.* 10 (1965)
2. Czarnuch, S., Mihailidis, A.: The design of intelligent in-home assistive technologies: Assessing the needs of older adults with dementia and their caregivers. *Gerontechnology* 10, 165–178 (2011)
3. Czarnuch, S., Mihailidis, A.: Depth image hand tracking from an overhead perspective using partially labeled, unbalanced data: Development and real-world testing (2014), <http://arxiv.org/abs/1409.2050v1>
4. Fontaine, J.R.J., Scherer, K.R., Roesch, E.B., Ellsworth, P.C.: The world of emotions is not two-dimensional. *Psychological Science* 18, 1050–1057 (2007)
5. Heise, D.R.: *Expressive Order: Confirming Sentiments in Social Actions*. Springer (2007)
6. Heise, D.R.: *Surveying Cultures: Discovering Shared Conceptions and Sentiments*. Wiley (2010)
7. Hoey, J., Boutilier, C., Poupart, P., Olivier, P., Monk, A., Mihailidis, A.: People, sensors, decisions: Customizable and adaptive technologies for assistance in healthcare. *ACM Trans. IIS* 2(4) (2012)
8. Hoey, J., Schröder, T., Alhothali, A.: Affect control processes: Intelligent affective interaction using a POMDP (2013), <http://arxiv.org/abs/1306.5279v2>
9. Hoey, J., Schröder, T., Alhothali, A.: Bayesian affect control theory. In: *Proc. Conf. on Affective Computing and Intelligent Interaction, ACII* (2013)
10. Lin, L.: *An Assistive Handwashing System with Emotional Intelligence*. Master's thesis, University of Waterloo (2014)
11. Lively, K.J., Smith, C.L.: Identity and illness. In: *Handbook of the Sociology of Health, Illness, and Healing*, pp. 505–525. Springer, New York (2011)
12. MacKinnon, N.J., Robinson, D.T.: 25 years of research in affect control theory. *Advances in Group Processing* 31 (2014)
13. Malhotra, A., Yu, C., Schröder, T., Hoey, J.: An exploratory study into the use of an emotionally aware cognitive assistant. TR# CS-2014-15, U. of Waterloo (2014)
14. Mihailidis, A., Boger, J., Candido, M., Hoey, J.: The coach prompting system to assist older adults with dementia through handwashing: An efficacy study. *BMC Geriatrics* 8(28) (2008)
15. Orona, C.J.: Temporality and identity loss due to Alzheimer's disease. *Social Science & Medicine* 30(11), 1247–1256 (1990)
16. Osgood, C.E., May, W.H., Miron, M.S.: *Cross-Cultural Universals of Affective Meaning*. University of Illinois Press (1975)
17. Rose Addis, D., Tippett, L.: Memory of myself: Autobiographical memory and identity in Alzheimer's disease. *Memory* 12(1), 56–74 (2004)
18. Scholl, W.: The socio-emotional basis of human interaction and communication: How we construct our social world. *Social Science Information* 52, 3–33 (2013)
19. Wilson, R., et al.: Quantitative analysis of formal caregivers' use of communication strategies while assisting individuals with moderate and severe alzheimer's disease during oral care. *Journal of Communication Disorders* 46(3), 249–263 (2013)