Decision theoretic, context aware safety assistance for persons who wander

Jesse Hoey[1], Xiao Yang[1], and Jesus Favela[2]

[1] School of Computer Science, University of Waterloo, Waterloo, Canada
[2] Computer Science Dept., CICESE, Ensenada, Mexico

Abstract. Wandering is one of the behaviors experienced by people with dementia that presents high risks and causes significant concern to caregivers. The frequency and manner in which a person wanders is highly influenced by the person's background and contextual factors specific to the situation. In this paper, we propose a decision-theoretic context-aware assistant that runs on a mobile (handheld) device, and uses multiple sources of local contextual information to provide verbal prompts, visual aids, or other help to the person when they wander. The system models uncertainty, learns about a person's patterns of behaviour, and can reason decision theoretically about the costs of sensors (e.g. battery charge) and the relative costs of different types of assistance. The system can be tailored to a particular person's needs by the user themselves or by their caregiver. The paper demonstrate a particular instance of this assistance system running on an Android platform.

1 Introduction

People with dementia suffer from spatial disorientation and memory loss, which causes them to get lost and wander, one of the behaviors of dementia least understood [1] and one of the main causes of concern among caregivers [9]. The Alzheimer's Association estimates that 60% of those suffering from dementia wander at some point, and half of those not found within a day suffer serious injuries or death. Solutions for the wandering problem usually take the form of a simple "virtual fence" that sends an alarm by text message or email to the caregiver when the user goes beyond a certain distance from home or provide a web interface for more detailed location information and monitoring [3]. However, these systems have a limited range of communication options, do not reason about long-term effects of actions or uncertainty, and do not take into account the wide variety of other contextual information that may be available, such as the detailed movements of the person, local network information, weather, battery state, etc. Further, sending alarms or notifications in situations where the user can find her way with a small bit of assistance will be an unnecessary burden on the caregiver, and infringes on privacy and independence. Current systems do not model these detailed preference tradeoffs. It is known that wandering encompasses a variety of behaviors [5] that are originated by diverse factors [1], and thus demand different types of interventions.

[3] see alz.org/comfortzone for example

Previous research has looked at this problem from a variety of angles. The Activity Compass [10] was a tool designed to help disoriented people find their destination. An extension of the Activity Compass project is Opportunity Knocks [7], a system designed to provide directional guidance to a user navigating through a city. iWander is a mobile application that uses contextual information such as the time of day and weather conditions to estimate the probability that the user of the mobile is wandering [12], but does not reason decision theoretically about the relative assistance mechanisms. Vuong et al. proposed an algorithm to detect wandering behaviors such as random, lapping, pacing and direct wandering [14].

In this paper, we discuss a novel context-aware sensing system, called LaCasa (Location and Context Aware Saftey Assistant), implemented on a handheld device that can reason about stochastic temporal events and make decision theoretic choices about help for a person with dementia (PwD). The system's controller is based on a partially observable Markov decision process, or POMDP [11]. One of the strengths of the POMDP formalism is that it allows designers to specify the model and a utility function, rather than a policy directly. The result is a more flexible design process that allows the system to be more easily customized for different users and different contexts. The POMDP model we describe in this paper was constructed in part based on our initial user requirements analysis and use-case scenarios [4].

2 Decision Theoretic Model

A POMDP consists of a finite set $\mathcal{S}$ of states; a finite set $\mathcal{A}$ 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 $\mathcal{S}$; a finite observation set $\mathcal{O}$; 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 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 $\mathcal{S}$) into choices of actions, such that the expected long-term discounted sum of rewards is (approximately) maximized.

The POMDP model is specified using the SNAP system [3], which is a user-friendly system for specifying POMDP models for assistance. The SNAP system breaks the state space down into three factors: *task* (T), *ability* (Y) and *behaviour* (B). The *task* variables are a characterisation of the domain. For example, in LaCasa, these include the location of the person and whether they are near a known location, along with additional context of the situation (e.g. time of day, weather, battery power). The task states are changed by the client's *behaviour*, B. In LaCasa, these include wandering or navigating to a known location. The client's *abilities* are their cognitive state, and model, e.g., the ability to recognise a known location, and the ability to find their way home.

The system actions are prompts that help the client regain a lost ability. There are two further actions: to *call_caregiver* (when the system is unable to help) and to (remind the person to) *recharge_battery* when the battery is running low. The observations are the sensor outputs, and measure the location (GPS or network), accelerations, their connectivity (wi-fi or cellular), and these give information about the person's current *task* state (e.g. are they wandering or not?). Complex virtual sensors can also be used, such as for detecting activities and social contact [13], affect [8] or location awareness (see Section 3). POMDPs can gracefully handle missing observations (from e.g. non-functioning sensors) by interpreting this as missing evidence in a Bayesian update.

The POMDP has an additonal variable for the battery level. This is critical to include explicitly in the model, because the POMDP can reason directly about the cost of querying different sensors, and can make decisions about which sensors to get information from in which states [6]. This can be an important consideration on a mobile device, where battery resources are at a premium. For example, the POMDP can reason about the measurement precision and battery cost trade-off for different location sensors (e.g. GPS and network-based).

3 Learning Patterns of Behaviour

A key element of LaCasa is to determine how likely it is that a person can *wayfind* (navigate to a known location) independently, possibly with assistance from LaCasa, but without assitance from a human. This allows the POMDP model to evaluate whether to call for human assistance, or to continue trying to help them. In the latter case, LaCasa may opt for different strategies based on the probability that the person can independently get home or to a known location. For example, they may be able to easily get to a coffee shop nearby with a small amount of assistance from LaCasa, but will not be able to make the longer trip all the way home, even with assistance. In some instnaces, just reaching a known location might reduce anxiety in the PwD, which in turn could make it easier for her to find her way home.

As LaCasa optimises battery life by only occasionally querying sensors, it only has a discrete set of points with labels showing whether the person was able to get home or not from there. Each time a person returns home, LaCasa saves the set of all locations $\mathbf{X} = \{\mathbf{x_1}, \mathbf{x_2}, \ldots, \mathbf{x_m}\} = \{x_1, y_1\}, \{x_2, y_2\}, \ldots, \{x_m, y_m\}$ from their most recent trajectory along with a label for each point $S_{\mathbf{x_i}} = 0$ if they required human assistance to get home and $S_{\mathbf{x_i}} = 1$ if they made it home independently. Note that the set of points $\mathbf{X}$ is at locations determined by LaCasa for battery efficiency reasons [6], not because of their geographic significance (e.g. distance apart in space or time).

LaCasa therefore has a set of N data points (from all trajectories so far) with labels $S_{\mathbf{x_1}} = s_1, S_{\mathbf{x_2}} = s_2, \ldots, S_{\mathbf{x_n}} = s_n$ as defined above, and needs to compute the probability that a person can make it home independently from any location $\mathbf{x_i} = \{x_i, y_i\}$. To accomplish this, we imagine a single Bernouilli experiment (weighted coin flip) at each location with an unknown probability $\theta(\mathbf{x}) = Pr(S_{\mathbf{x}} = 1)$. We then use a beta distribution as a conjugate prior for each experiment, and we make the assumption that locations close to one another (with distance measured by a kernel function $K(\mathbf{x}, \mathbf{x'})$) will have similar probabilities. After observing the outcome $\mathbf{s}$ of a set of experiment $S_{\mathbf{x_1}}, \ldots, S_{\mathbf{x_n}}$, we update the distribution over $\theta(\mathbf{x})$ to obtain the approximate posterior [2]:

$$\begin{aligned} &\Pr(\theta(\mathbf{x}) | S_{\mathbf{x_1}} = s_1, S_{\mathbf{x_2}} = s_2, \ldots, S_{\mathbf{x_n}} = s_n) \\ &\quad \propto \theta(\mathbf{x})^{\alpha(\mathbf{x})-1+\sum_{i=1}^{n} \delta(s_i=1) K(\mathbf{x_i}, \mathbf{x})} \times (1-\theta(\mathbf{x}))^{\beta(\mathbf{x})-1+\sum_{i=1}^{n} \delta(s_i=0) K(\mathbf{x_i}, \mathbf{x})} \end{aligned} \tag{1}$$

where $\delta(a)$ is a Dirac delta that returns 1 when a is true and 0 otherwise. The resulting posterior $\Pr(\theta(\mathbf{x}))$ is an estimate of the distribution of the probability that the person will wayfind independently. We can then use the mean of this distribution $\mu_\theta = \alpha/(\alpha+\beta)$, as a point estimate of the probability, or we can use a confidence adjusted version $\mu_\theta - \gamma\sigma_\theta$, where $\sigma_\theta = \sqrt{\alpha\beta}/((\alpha+\beta)\sqrt{\alpha+\beta+1})$ is the standard deviation and ν is a constant between -1 and 1 that ranges from optimistic to pessimistic estimates. This probability map will be used as a virtual sensor giving indications of whether a person is in a known location or not (see Section 4). Note that conflicts in labels are handled gracefully, the probabilities adjusting based on the relative frequency of the two labels.

We have considerable flexibility in choosing the kernel function, K, and here we consider that, if a person can wayfind from location $\mathbf{x}$, they will be more likely to wayfind from another location $\mathbf{x}'$ if $\mathbf{x}$ and $\mathbf{x}'$ are (i) close in walking distance (not as the crow flies) and (ii) within sight. We therefore use $K(\mathbf{x},\mathbf{x}') = e^{-(c(\mathbf{x},\mathbf{x})d(\mathbf{x},\mathbf{x})/b)^2}$, where $d(\mathbf{x},\mathbf{x})$ is the walking distance, b is the kernel bandwidth, and $c(\mathbf{x},\mathbf{x}')$ is a *line-of-sight* function that is < 1 and smaller if $\mathbf{x}$ and $\mathbf{x}'$ lie within sight.

Figure 1 shows three examples given a set of randomly generated points at intersections in a latitude-aligned grid of streets in which the PwD can more easily wayfind if they are closer to the central value of y. We use a *line-of-sight* function $c = C$ if the locations x and x' lie on the same latitude or longitude (but not both) to within the street width, and $c = 1$ otherwise. Figure 1 shows the learned probability functions for three settings of the parameter C. As C gets smaller, the *line-of-sight* is more meaningful, and the effect of a person wayfinding from a certain location "spreads" along the associated streets.

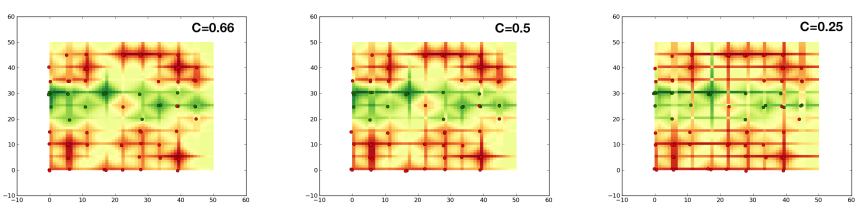

Fig. 1. Examples of learned probability map for a person for different C values.

4 Mobile Implementation

The system is implemented on Google's Android (Gingerbread) operating system. The POMDP belief updates and policy queries are handled by a remote desktop machine that communicates with the smartphone over TCP/IP using simple XML messages. The device, on starting, queries the server and gets information about the set of observations and actions that its POMDP model can handle. LaCasa runs as a foreground service on the Android platform, and has a list of sensors (on the device or remotely in the local environment) that it is currently able to get information from. It can register listeners for each of these sensors, and each sensor has a method to convert the raw sensor readings into a discrete observation label for a particular POMDP observation variable. Adding a new sensor involves implementing an abstract sensor class, defining sensor conversion methods, and publishing the new sensor name. When a user selects a

POMDP model to use (from a set they have created using SNAP), the LaCasa service queries the POMDP server for information about this model, and registers a new process with the server for the model. Subsequently, upon receiving an updated sensor reading, the Android device updates its current observation set and sends the subset required by the model to the server. The POMDP running on the server updates its belief and consults its policy, returning an action from a discrete set to the mobile device. The mobile then has a set of user-defined action mappings that translate this discrete action label into an action to take.

	Observations			Task						Behaviour					Ability						
Step	known_loc	location	at_known_location	at_home_no	at_home_yes	known_location_close_no	known_location_close_yes	in_known_location_no	in_known_location_yes	behaviour_other	behaviour_nothing	behaviour_wander	behaviour_go_home	behaviour_go_to_known	know_path_to_known_yes	know_path_to_known_no	rn_known_location_yes	rn_known_location_no	know_path_home_yes	know_path_home_no	Action for System
0	close	home	yes																		donothing
3	close	far_from_home	yes																		donothing
4	close	home	yes																		know_path_home
7	close	close_to_home	no																		call_caregiver
8	far	far_from_home	no																		rn_known_location
9	close	far_from_home	no																		know_path_to_known
10	close	far_from_home	no																		rn_known_location
18	far	far_from_home	no																		know_path_to_known
25	close	far_from_home	no																		rn_known_location
32	close	home	yes																		donothing

Table 1. Example simulation for LaCasa, battery and caregiver variables not shown.

We now present a simple example of a LaCasa system that has a variable describing whether the persons is *at-home* or not, and another describing whether they are close to a *known_location*, and a third describing whether they are at a known location currently. Known locations are those from which the user can more easily find their way home, e.g. known landmarks. A person will be deemed to be "at" a known location if the device can connect to a trusted wireless network (wi-fi), or if the $P(\theta(x))$ (Section 3) is above a threshold (these two evidence sources have different strengths, wireless connectivity being the stronger). The modeled abilities are whether the PwD knows the path home, knows the path to a known location, and recognizes a known location (*rn_known_location*). The observations are the persons *location* $\in$ {*home*, *close_to_home*,*far_from_home*}, the battery charge, whether they are at a known location, whether they are close to a known location, and whether the caregiver is present or not.

Table 1 shows part results of a demonstration carried out with the mobile device. A test subject (one of the authors) started at their home location, and went out for a walk. There are four known locations in the vicinity of their home, but none of them have a trusted wi-fi connection, so they are simply locations the person might know how to get to, rather than locations where the person can spend some time. Initially, when in close proximity of their home, the system does nothing and only monitors the user. When the person moves far enough from home to have their location labeled as *far-from-home*, at which point the system prompts them for the path homewards (step 4). The person ignores that prompt, and continues to move farther from the home. The system then calls for caregiver assistance (step 7), but continues to try to help the person find a known location (steps 8-25). Once the person regains their home, the system returns to only monitoring (step 32).

5 Conclusions and Future Work

Wandering is a common and complex behavior among PwDs while being one of the major causes of concern among caregivers. We have proposed the use of a partially observable Markov decision process to account for the many forms in which this behavior is expressed and the different actions that are required to deal with it. In this paper, we further propose a framework for learning the behaviours of a client based on data. The model has been implemented on a mobile device that is used to gather contextual information to feed the model, to provide assistance to the PwD and to communicate with the caregiver when considered appropriate. Our future aims are to further develop the learning methodology, and to elicit requirements from end users to further define the system. This includes tackling the problem of how to design effective prompts using mobile device notifications that do not increase user confusion.

References

1. Algase, D.: Wandering in dementia. Ann. Rev. Nursing Res. 17, 189–218 (1999)
2. Goetschalckx, R., Poupart, P., Hoey, J.: Continuous correlated beta processes. In: Proc. Intl. Joint Conf. on Artificial Intelligence (IJCAI). (July 2011)
3. Hoey, J., Plötz, T., Jackson, D., Monk, A., Pham, C., Olivier, P.: Rapid specification and automated generation of prompting systems to assist people with dementia. Pervasive and Mobile Computing 7(3), 299–318 (June 2011)
4. Hoey, J., Yang, X., Quintana, E., Favela, J.: LaCasa:location and context aware safety assistant. In: Proc. of 6th International Conference on Pervasive Computing Technologies for Healthcare. San Diego, CA (2012)
5. Hope, R., Fairburn, C.: The nature of wandering in dementia: a community based study. Int J Geriatric Psychiatry 5, 239–245 (1990)
6. Koltunova, V., Hoey, J., Grześ, M.: Goal-oriented sensor selection for intelligent phones: (GOSSIP). In: Proc. SAGAware '11, ACM, (2011),
7. Liao, L., Patterson, D., Fox, D., Kautz, H.: Learning and inferring transporation routines. Artificial Intelligence 171(5-6), 311–331 (2007)
8. Matic, A., Papliatseyeu, A., Gabrielli, S., Osmani, V., Mayora-Ibarra, O.: Happy or moody? why so? monitoring daily routines at work and inferring their influence on mood. In: Proc. Ubi-Health (2010)
9. Miyamoto, Y., Ito, H., Otsuka, T., Kurita, H.: Caregiver burden in mobile and non-mobile demented patients. Intl. J. Ger. Psych.17(8), 765–773 (2002),
10. Patterson, D., Etzioni, O., Kautz, H.: The activity compass. In: First International Workshop on Ubiquitous Computing for Cognitive Aids (Sept 2002)
11. Poupart, P.: An introduction to fully and partially observable Markov decision processes. In: Sucar, E., Morales, E., hoey, J. (eds.) Decision Theory Models for Applications in Artificial Intelligence, chap. 3, pp. 1–30. IGI Global (2011)
12. Sposaro, F., Danielson, J., Tyson, G.: iwander: An android application for dementia patients. In: Conf Proc IEEE Eng Med Biol Soc. pp. 3875–8 (2010)
13. Tung, J.Y., Semple, J.F., Woo, W.X., Hsu, W.S., Sinn, M., Roy, E.A., Poupart, P.: VALMA: Voice, activity and location monitoring for alzheimer's disease and related dementias. In: RESNA, Toronto (2011)
14. Vuong, N.e.a.: Feasibility study of a real-time wandering detection algorithm for dementia patients. In: MobileHealth '11: Proceedings of the First ACM MobiHoc Workshop on Pervasive Wireless Healthcare (May 2011)