

---

# Modeling and Learning for LaCasa, the Location And Context-Aware Safety Assistant

---

**Jesse Hoey, Xiao Yang, Marek Grzes**

School of Computer Science

University of Waterloo

Waterloo, ON, N2L 1Y4

{jhoey, x85yang, mgrzes}@uwaterloo.ca

**Rene Navarro, Jesus Favela**

CICESE

Ensenada, Mexico

{rnavarro, favela}@cicese.mx

## Abstract

Wandering is a common behavior among people with dementia (PwD). It is also one of the main concerns of caregivers since it can cause the person to get lost and injured. 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 investigate some of the properties of wandering behaviours, particularly related to our ability to sense them with mobile devices. We then present a novel decision-theoretic model that estimates the risk faced by the PwD and decides on the appropriate action to take, such as prompting the PwD or calling the caregiver. The model can be tailored to the user needs (e.g. known locations, level of cognitive decline) takes into account uncertainty, seamlessly merges data from different sensors, reasons about the relative worth of sensor readings (and so can be used for active sensing), and learns about a person's patterns of behaviour. Contextual information gathered from sensors is seamlessly factored into the model, including current location, battery power, and proximity to the caregiver. The system can reason decision theoretically about the costs of sensors (e.g. battery charge) and the relative costs of different types of assistance. A preliminary version of the system has been instantiated in a wandering assistance application for mobile devices 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 [2, 3]. 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. Location-aware technology has been proposed to assist caregivers with the risks associated with wandering, however, wandering encompasses a variety of behaviors [4] that are originated by diverse factors [1], and thus demand different types of interventions.

Wandering behavior is unique to each person in each context or situation [5]. The factors that affect wandering include those related to the background of the person such as health status, personality, cognitive decline, and socio-demographic characteristics, and proximal factors related to the context of the situation including personal and physical needs as well as physical and social aspects of the environment [5]. For example, deficits in visual perception (background factor) reduce cues that aid in interpreting and navigating an environment with deficient lighting (proximal factor).

Understanding these differences is paramount to the development of technological solutions aimed at assisting persons with dementia (PwDs) and their caregivers. In this paper we describe a solution that incorporates some of these background and contextual factors to detect risky wandering

behavior, and to decide on the type of assistance to be provided [6, 7]. The system is based on a partially observable Markov decision process, or POMDP, a stochastic and decision theoretic model of a temporal dynamic process that encodes contexts and preferences in a general purpose and mathematically well founded way. The POMDP can be used to generate a *policy* for the assistance system that is well defined and interpretable, and is easily modified to suit different user types and contexts.

Section 2 presents the typology and causes of wandering, illustrated with two scenarios. Section 4 describes how contextual information relevant to wandering is obtained from mobility data from older adults. Sections 3 presents the decision theoretic model, Section 5 overviews our work on learning, and 6 shows a prototype implementation. Sections 7 and 8 present related work and conclusions.

## 2 Characterizing Wandering Behaviors

Wandering is defined by the North American Nursing Diagnosis Association (NANDA) as "meandering, aimless, or repetitive locomotion that exposes a person to harm and frequently is incongruent with boundaries, limits, or obstacles" [8]. It is estimated that more than 60% of dementia sufferers wander, yet its causes and manifestations are varied.

Wandering behaviors have been classified as random walking (ambulation having frequent direction changes), lapping (moving in circuitous or looping fashion), pacing (moving back and forth repetitively between two points) and direct ambulation [9]. Random wandering becomes more prevalent with increased cognitive impairment while pacing is less frequent, is not significantly associated with cognitive impairment, and might be more an indication of agitation or anxiety. Of particular concern is direct or goal-oriented wandering, which occurs at all levels of cognitive impairment and usually responds to physiological or psychological needs that could be anticipated. Direct wandering often involves longer distances and leads to the PwD getting lost as he tries to reach the house where he used to live or the place where he used to work.

We use scenarios based on actual cases to illustrate some of the factors influencing wandering behavior and the different types of support required by a PwD who wanders, and to inform the design of the support tool. These differences highlight the need to incorporate background information about the individual and contextual information about the situation.

*Scenario 1. Getting lost from spatial disorientation.* Marie is a 70-year-old woman who lives in a small village in France. She was recently diagnosed with Alzheimer's disease, and has increasing trouble remembering where she placed things and what her schedule is. She lives with her husband, who is also 70 but does not have a memory impairment. One of Marie's greatest pleasures is to take her dog for a walk, which helps her deal with the increased anxiety and depression that affects her as her disease progresses. At times she feels disoriented, and has been brought back home a few times by other villagers. She feels increasingly inhibited to go out on walks because of the social stigma attached to this dependence, and her husband feels increasingly anxious when she leaves. She has a number of friends in the village, and regularly frequents a few of the shops, from which she can easily get home.

*Scenario 2. Getting lost due to goal-oriented wandering.* Antonio worked for more than 40 years in a vineyard. He retired in large part because he was diagnosed with Alzheimer's disease five years ago. Often Antonio becomes anxious and tells his wife that he needs to go to work, and she has increasing difficulty convincing him that he has already retired. On a few occasions she has found him a block or two away from home walking in the direction to his former workplace and with some effort she manages to convince him to get back home or walks with him until he forgets where he was going. Once, he took the bus late in the afternoon to the vineyard and arrived there feeling disoriented. Fortunately, a former colleague of his found him and notified his wife.

In the first scenario, Marie is mostly aware that she is lost, she gets disoriented due to lack of attention and poor visual and spatial ability, even though she might be in an area she is familiar with. In her case, the anxiety caused by feeling lost can be lessened by providing her with anchors to reality, e.g., reminding her that she was heading home or showing her a close-by known location, such as her local bakery.

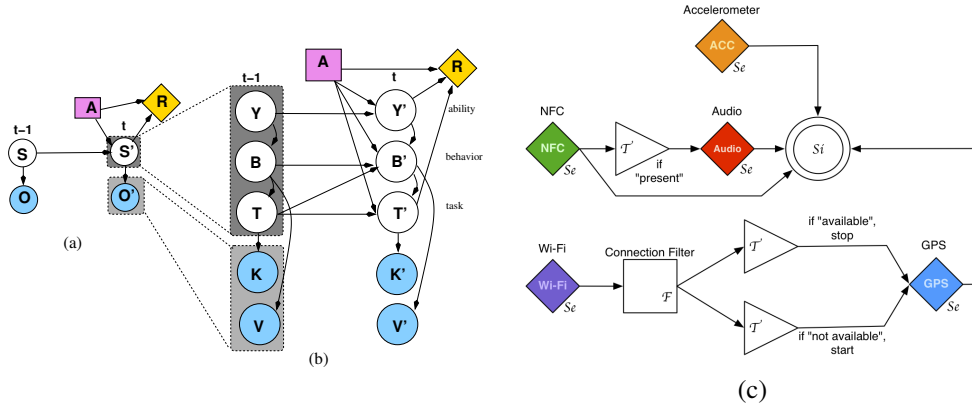


Figure 1: Two time slices of (a) a general POMDP; (b) a factored POMDP for modelling assistive technology; (c) Model depicting a sensing session using InCense.

The second scenario is an instance of direct wandering, as Antonio does not feel lost when he starts to walk to his former workplace. On the contrary, his wandering is a manifestation of his need to go to work and he will not be receptive to suggestions to return home or ask for assistance, so automatic prompting might lead to more confusion and anxiety. In this case notifying a caregiver or someone in the vicinity who can provide assistance would be more effective.

These two scenarios also show how independence and safety need to be balanced according to the profile of the PwD and the factors that lead to wandering. Antonio, who is at a more advanced stage of dementia has a greater risk of getting lost and might travel larger distances. While he might have unmet needs that cause the wandering, independence is not a major issue in his quality of life as he is already under the constant care and supervision of his wife. In contrast, for Marie, being able to walk her dog along the beach gives her satisfaction and contributes to her fitness. The risks she faces of getting disoriented in the small village where she is well known do not justify her refraining from going outdoors on her own.

### 3 Decision Theoretic Model

In this section, we describe a context-aware sensing system called LaCasa, that is implemented on a handheld device and that can reason about stochastic temporal events and make decision theoretic choices about help to offer [6, 7]. This assistance can take the form of audio or visual prompts, explicit guidance to close-by known locations, or SMS messages. The system's controller is a partially observable Markov decision process, or POMDP [10], a stochastic model of a temporal dynamic process that can handle partial observability (i.e. elements of the state that are not directly measurable, but only inferrable from some sensor readings). The model includes a set of preferences encoded as utilities, and computes a policy of assistance that optimises over the utilities given the state and dynamics of the world in the long term. Sensor accuracies are naturally folded into the decisions. Assistance systems for persons with dementia for activities of daily living (ADL) such as handwashing have been built using POMDPs [11].

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  $\Pr(t|s, a)$  denoting the probability of moving from state  $s$  to  $t$  when action  $a$  is taken, and  $\Delta(\mathcal{S})$  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  $a$ . A generic POMDP is shown as a Bayesian network in Figure 1(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 monitor beliefs about the system state using standard Bayesian tracking/filtering. A *policy* can be computed that maps these *belief states* into choices of actions, such that the expected discounted sum of rewards is (approximately) maximized. We compute policies using a point-based value iteration approach (see [10] for details and an introduction to POMDPs). POMDPs are a very general purpose model for systems that interact autonomously with an environment: they can be adapted to different contexts, sensors and actions depending on user requirements. Specification of a POMDP involves describing the domain in which the system is to operate, and what the goals of the system are. This is in contrast with the more traditional approach in which a deterministic controller is designed (a hand-crafted or rule-based policy). Although it may be possible to engineer a successful rule-based system, it will be less flexible and extensible as user needs change or as novel contexts arise. The POMDP, on the other hand, uses a consistent and well-defined decision theoretic model in which user needs are clearly defined using utility theory, and contexts can be easily specified using a task analysis tools such as SNAP [12] (as we use here). For example, our two scenarios can be distinguished in the POMDP simply by increasing the cost of prompting in Antonio’s case (as it should be used with more care), or increasing the cost of calling a caregiver in Maria’s case (as she finds this intrusive). These increases in cost have a well-defined meaning, and can be elicited from users using a variety of standard *preference elicitation* techniques.

In order to specify an initial version of the POMDP controller, we have developed a user-friendly system for specifying POMDP models for assistance called SNAP [12]. The SNAP task analysis system breaks the state space down into three factors based on a psychological “syndetic” model that captures the important aspects of these types of domains. These three factors are shown in Figure 1(b). The state,  $S$ , has been factored into three sets of variables: *task* ( $T$ ), *ability* ( $Y$ ) and *behaviour* ( $B$ ). The *task* variables are a characterisation of the domain in terms of a set of high-level variables. 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, 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 or to home. Doing nothing or something unrelated are always included. The client’s *abilities* are their cognitive state, and model, e.g., the ability to recognise a known location or destination, and the ability to find their way home.

Notice that the task state,  $T'$ , is independent of the ability,  $Y'$ , given the behaviour,  $B'$ . Changes in the *task* states are caused by the *behaviours* of the client (and possibly the system actions), independently of the client’s ability given the behaviours. However, the *behaviours* are directly influenced by the *abilities* (so a person who can recall what to do is more likely to do it).

The system actions are prompts that help the client regain a lost ability. We define one system action for each necessary ability in the task. The actions correspond to a prompt or signal that will help the client with this particular ability, if missing. For example, the mobile platform can sense the position of the person (using GPS or network locators), 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?). More complex virtual sensors can also be used, such as activity recognition systems based on acceleration or voice data [13], which can give evidence about the *behaviours* or *abilities* of the person. For example, goal oriented wandering could be detected by a virtual sensor that measures their progress towards a particular location over an interval of time. Disoriented wandering could be measured by looking for trajectories that are not optimal [14].

Finally, the utility function in the POMDP indicates which states are desirable. These may change from user to user, but may include factors such as their desire for independence, or the cost of calling a caregiver.

LaCasa is encoded in SNAP, and includes 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 can be locations that the person can easily recognise such as landmarks or restaurants. A person will be deemed to be “at” a known location if the device can connect to a trusted wireless network (wi-fi). Their abilities are to know the path home, know the path to a known location, and to recognize a known location (*rm\_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. However, this obser-

vation set can be expanded easily without changing the underlying POMDP model. Similarly, the set of relevant contextual variables can be adjusted for each person using SNAP.

## 4 Automatic estimation of wandering context

A field study was conducted to assess the feasibility of obtaining accurate information to be used by the decision theoretic model proposed in this paper. In particular we aimed at detecting locations frequently visited by the older adults to obtain base data from which to infer anomalous behavior. Behavioral data from 15 community-dwelling older adults was gathered during 20 days using smartphones. Five men and ten women over 70 (average age 75) participated in the study, three of the participants lived alone, six only with their spouses and six with other family members.

Participants were asked to carry the smartphone on their hip at all times except when they went to bed or took a bath. Data was gathered using the InCense application, a toolkit for behavioral data collection from populations of mobile phone users that supports an opportunistic and/or participatory sensing paradigm [15]. InCense provides a GUI used by the researcher to create a sensing campaign using mobile devices running Android. A data gathering session created with InCense includes sensors, filters, triggers and data sinks. Figure 1(c) shows part of the model created to gather the data used in this project. It shows, e.g., that to conserve battery life, GPS location was gathered only when InCense fails to detect an access point at the user’s home, where the participants spent on average 21.5 hrs/day. To gather ground truth, participants were asked to place the cellular phone on top of NFC tags before initiating an activity. The tags, which include pictograms of relevant activities, were distributed throughout the house in locations associated with the activity.

We conducted a preliminary analysis of the data to automatically determine locations frequented by the participants. This information is used by the decision theoretic model to determine if the PwD is lost and to offer him navigational assistance. The list of known locations is determined by selecting spatial coordinates where at least two consecutive recordings fall within a diameter of 50 meters and in which the user stays for at least 6 minutes. When a candidate location falls within the area of another known location it is eliminated. Following this approach we obtained only one false positive when we compared with results obtained through visual inspection and self-report. There were on average 2 false negatives per participant from an average of 16 frequented locations per participant. These are mostly locations where the person spends very little time, and include gas stations where they stop to fill their car tank, a school where an elder drops and picks-up their grandchildren, and bus stops.

## 5 Learning Models and Behaviour Patterns

It is possible to learn the POMDP models from the data we gather over time by gathering statistics on the belief states visited, observations, actions and rewards. For example, we can learn the proximal factors that trigger wandering, such as time of day, weather, socialization, previous activity intensity, etc. In one of our scenarios, a PwD often experiences “sundowning” syndrome in the evening and goes to her room, starts putting clothes in her suitcase with the intention of going to the town in which she used to live. The POMDP can learn that the dimming light is one of the main triggers of this behavior. Other types of wandering, such as pacing, are often triggered by the need to move the body after a period of rest (e.g., a nap), and could be picked up by an accelerometer. Another example would be “discussion with the caregiver” giving rise to anxiety leading to wandering and withdrawal. Shouting and intense discussions could potentially be obtained from audio analysis. An analysis of this data could help a social worker or physician analyze with the caregiver potential causes and strategies to ameliorate this behavior. Thus, the POMDP can also be used as an assessment tool.

A second key area of learning is the determination of how likely it is that a person can *wayfind* (navigate to a known location) independently, possibly with assistance from LaCasa, but without assistance 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.




Figure 2: Examples of learned probability map for a person for different C values.

In some instances, 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  $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m\} = \{x_1, y_1\}, \{x_2, y_2\}, \dots, \{x_m, y_m\}$  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 [16], 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  $S_{\mathbf{x}_1} = s_1, S_{\mathbf{x}_2} = s_2, \dots, S_{\mathbf{x}_n} = s_n$  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 Bernoulli 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  $S_{\mathbf{x}_1}, \dots, S_{\mathbf{x}_n}$ , we update the distribution over  $\theta(\mathbf{x})$  to obtain the approximate posterior [17]:

$$\begin{aligned} & \Pr(\theta(\mathbf{x}) | S_{\mathbf{x}_1} = s_1, S_{\mathbf{x}_2} = s_2, \dots, S_{\mathbf{x}_n} = s_n) \\ & \propto \theta(\mathbf{x})^{\alpha(\mathbf{x})-1 + \sum_{i=1}^n \delta(s_i=1)K(\mathbf{x}_i, \mathbf{x})} \times (1 - \theta(\mathbf{x}))^{\beta(\mathbf{x})-1 + \sum_{i=1}^n \delta(s_i=0)K(\mathbf{x}_i, \mathbf{x})} \end{aligned} \quad (1)$$

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  $\nu$  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 6). 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 2 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 2 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.

## 6 Mobile Implementation

The system is implemented on Google’s Android (Gingerbread) operating system (see screenshots in Figure 3). 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. Otherwise, the portable device only needs a data (internet) connection for its (limited) communication with the server. It is possible to push the POMDP controller onto the smartphone, either directly (belief monitoring and policy evaluation) or indirectly by compiling the policy into a finite state controller (FSC). The FSC option can be used to reduce battery consumption and relieve the user of the reliance on a data connection, but the precise tradeoffs are currently under investigation.

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, LaCasa registers a new process with the POMDP server, and requests information about the model. Subsequently, upon receiving an updated sensor reading, the device sends the subset required by the model to the server. The POMDP on the server updates its belief and consults its policy, returning an action label that the device translates into an action to take using a set of user-defined mappings. The user can change to a different POMDP (for some other context or purpose) by reconnecting to the server with a new model name. As we remarked in the previous section, a caregiver can customise a new POMDP model for a particular person’s needs, and the mobile system can easily make use of this without any changes on the mobile device.

For the wandering application, the person’s location (latitude and longitude) is converted into a set of discrete locations. In our simple model, we use three possibilities: *home* means she is within sight, *close-to-home* means she can find her way back easily, and *far-from-home* means she may have more trouble finding her way back home. We allow users to specify these regions manually by drawing on a map (Figure 3, middle), or we learn them using either heuristic methods as described in Section 4, statistical clustering methods, or Bayesian methods based on the correlated Beta process (Section 5 and [6]). To detect if a person is at a known location, we use the device’s *wi-fi* connectivity: if they are connected to a known wi-fi network, we can assume they are likely to be in a known location. LaCasa keeps a database of known locations (latitude/longitude) with one special home location, and compares current location estimates to the database. Users can also have special locations for old workplaces to aid with detecting goal-oriented wandering. Users can add/remove known locations and can change their home location. A photograph of each known location is stored, and can be displayed to the user as a possible prompt. The system can use its built-in text-to-speech system to prompt, can display photos of known locations, or can send SMS messages. It is also possible to have the POMDP model be aware of the cost of querying sensors, and be able to make decisions about which sensors to get information from in which states [16]. This can be an important consideration on a mobile device, where battery resources are at a premium. For example, querying a GPS sensor for location information is slow and expensive, but gives accurate results. On the other hand, location can be obtained from the network (cellular towers and wi-fi connections), and this is fast and cheap, but gives less precise results [18]. The POMDP model can then reason about this trade-off given the current user situation and query the appropriate sensor.

## 7 Related Work

Commercial technological solutions to wandering usually take the form of a simple “virtual fence” that sends an alarm to the caregiver when the PwD goes beyond a certain distance from home or provide a web interface for more detailed monitoring. Although this approach is practical and appropriate for many users, it fails to account for the variety of factors that lead to wandering and the different ways in which both the caregivers and PwD wish to deal with it. For instance, these solutions put safety over independence and privacy which might not be appropriate for users like Marie in our first scenario. Indeed in a focus group conducted with PwDs to assess the effectiveness of non-pharmacological interventions to wandering, all 6 participants stressed the importance of maintaining the independence and the need for autonomy [19].

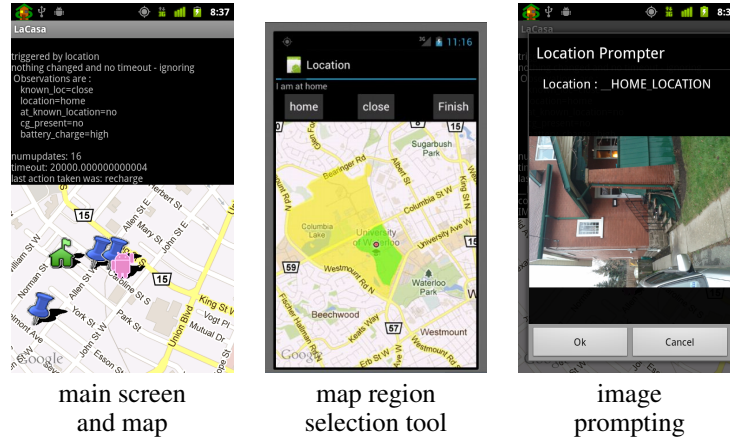


Figure 3: Example screen capture images from LaCasa prototype

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 [20]. A Bayesian network is used to assess the user’s risk and a message can be sent to the caregiver. iWander provides more flexibility than the “virtual fence” approach by taking context into consideration, but offers little flexibility in the type of assistance, being limited to either giving guidance or alerting a caregiver.

Vuong et al. proposed an algorithm to detect wandering behaviors such as random, lapping, pacing and direct wandering [14]. Temporal and spatial information is used to detect lapping and pacing, while random walking is inferred when the person follows a non-optimal route. All other walking behaviors are considered direct wandering, which fails to distinguish between being lost while trying to reach a destination, as with Antonio in Scenario 2, from safely walking to a known location. Liao *et al.* learned patterns of motion outside, including transportation routines, and were able to make accurate predictions about the likely goal locations of persons[21], specifically with a focus on providing assistance to the elderly. Their models would fit well into the control mechanisms that we have described in this paper.

## 8 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. The preliminary model takes into consideration the background of the PwD and information related to the context of the situation, to both detect a risky wandering episode and act accordingly. The model has been implemented on a mobile device that is used to gather contextual information, to provide assistance to the PwD and to communicate with the caregiver when considered appropriate. This is the first demonstration of a POMDP controller for assistance with dementia running on a mobile platform.

For the moment, the prototype incorporates only a few contextual variables. We will work on identifying those that are more relevant to accurately infer relevant behaviors and as appropriate actions to be taken by the system, to gradually incorporate them into a richer and more flexible model. Contextual factors would then condition our estimates that a person can effectively wayfind (e.g. recognise a known location). We could also model the persons overall state of mind on each day - so if they are having a lot of trouble one day (doing other things, like washing hands, dressing, etc), then we would lower our estimate that they could recognise a location. The interface between the device and the user is clearly a key challenge, and we plan to continue work on our interactive system that allows users to tailor the model to their specific needs [22]. We are also studying the different methods for implementing the POMDP policy on the mobile device. As battery power is at a premium, we are comparing our server-based policy method with one where the policy is converted to a finite state controller and pushed onto the phone. The tradeoffs between accuracy and cost are being studied.



**Acknowledgements:** This work was funded by the American Alzheimer’s Association grant number ETAC-10-173237.

## References

- [1] D. Algase, “Wandering in dementia,” *Annual Review of Nursing Research*, vol. 17, pp. 189–218, 1999.
- [2] Y. Miyamoto, H. Ito, T. Otsuka, and H. Kurita, “Caregiver burden in mobile and non-mobile demented patients: a comparative study,” *International Journal of Geriatric Psychiatry*, vol. 17, no. 8, pp. 765–773, 2002. [Online]. Available: <http://dx.doi.org/10.1002/gps.694>
- [3] A. Nelson and D. Algase, *Evidence-based protocols for managing wandering behaviors*. Springer, 2007.
- [4] R. Hope and C. Fairburn, “The nature of wandering in dementia: a community based study,” *Int J Geriatric Psychiatry*, vol. 5, pp. 239–245, 1990.
- [5] D. Algase, C. Beel-Bates, and E. Beattie, “Wandering in long-term care.” *Annals of Long-Term Care: Clinical Care and Ageing 2003*, vol. 11, no. 1, pp. 33–39, 2003.
- [6] J. Hoey, X. Yang, and J. Favela, “Decision theoretic, context aware safety assistance for persons who wander,” in *Proc. of 6th International Workshop on Ubiquitous Health and Wellness (UbiHealth)*, Newcastle, UK, 2012.
- [7] J. Hoey, X. Yang, E. Quintana, and J. Favela, “LaCasa: location and context aware safety assistant,” in *Proc. of 6th Intl. Conf. Pervasive Computing Technologies for Healthcare*, San Diego, CA, 2012.
- [8] NANDA, *NANDA, Nursing Diagnosis: Definitions and Classifications 2001-02*, Philadelphia, PA.
- [9] D. Martino-Saltzman, B. B. Blasch, R. D. Morris, and L. W. McNeal, “Travel behavior of nursing home residents perceived as wanderers and nonwanderers,” *The Gerontologist*, vol. 31, no. 5, pp. 666–672, 1991.
- [10] P. Poupart, “An introduction to fully and partially observable Markov decision processes,” in *Decision Theory Models for Applications in Artificial Intelligence: Concepts and Solutions*, E. Sucar, E. Morales, and J. Hoey, Eds. IGI Global, 2011, ch. 3, pp. 1–30.
- [11] J. Hoey, P. Poupart, A. von Bertoldi, T. Craig, C. Boutilier, and A. Mihailidis, “Automated handwashing assistance for persons with dementia using video and a partially observable Markov decision process,” *Computer Vision and Image Understanding*, vol. 114, no. 5, pp. 503–519, May 2010.
- [12] J. Hoey, T. Plötz, D. Jackson, A. Monk, C. Pham, and P. Olivier, “Rapid specification and automated generation of prompting systems to assist people with dementia,” *Pervasive and Mobile Computing*, vol. 7, no. 3, pp. 299–318, June 2011.
- [13] J. Y. Tung, J. F. Semple, W. X. Woo, W.-S. Hsu, M. Sinn, E. A. Roy, and P. Poupart, “Valma: Voice, activity and location monitoring for alzheimer’s disease and related dementias,” in *Annual Conf. of the Rehabilitation Engineering and Assistive Technology Society of North America (RESNA)*, Toronto, 2011.
- [14] N.K Vuong et al, “Feasibility study of a real-time wandering detection algorithm for dementia patients,” in *Proc. of the First ACM MobiHoc Workshop on Pervasive Wireless Healthcare*, May 2011.
- [15] M. Perez, L. Castro, and J. Favela, “Incense: A research kit to facilitate behavioral data gathering from populations of mobile phone users,” in *Proceedings of the 5th international symposium of ubiquitous computing and ambient intelligence*, ser. UCAMI’11, 2011.
- [16] V. Koltunova, J. Hoey, and M. Grześ, “Goal-oriented sensor selection for intelligent phones: (gossip),” in *Proceedings of the 2011 international workshop on Situation activity and goal awareness*, ser. SAGAware ’11. New York, NY, USA: ACM, 2011, pp. 83–88.
- [17] R. Goetschalckx, P. Poupart, and J. Hoey, “Continuous correlated beta processes,” in *Proc. International Joint Conferences on Artificial Intelligence (IJCAI)*, Barcelona, Spain, July 2011.
- [18] A. Carroll and G. Heiser, “An analysis of power consumption in a smartphone,” in *Proceedings of USENIX Annual Technical Conference*, Boston, MA, June 2010.
- [19] L. Robinson, D. Hutchings, L. Corner, F. Beyer, H. Dickinson, A. V. A, T. Finch, J. Hughes, C. Ballard, C. May, and J. Bond, “A systematic literature review of the effectiveness of non-pharmacological interventions to prevent wandering in dementia and evaluation of the ethical implications and acceptability of their use.” *Health Technology Assessment*, v. 10, n. 26, 2006.
- [20] F. Sposaro, J. Danielson, and G. Tyson, “iwander: An android application for dementia patients,” in *Conf Proc IEEE Eng Med Biol Soc.*, 2010, pp. 3875–8.
- [21] L. Liao, D. Patterson, D. Fox, and H. Kautz, “Learning and inferring transportation routines,” *Artificial Intelligence*, vol. 171, no. 5-6, pp. 311–331, 2007.
- [22] A. Hwang and J. Hoey, “Smart home, the next generation: Closing the gap between users and technology,” in *AAAI Fall Symposium on Gerontechnology*. Arlington, VA: AAAI, 2012.