Home Sensor Data Fusion to Support Aging in Place

FINAL REPORT

12 APRIL 2005

Sponsored by

National Institutes of Health – National Institute on Aging
Small Business Innovation Research Grant

1 R43 AG024687-01

Submitted to:
Agency for Healthcare Research and Quality
Grants Management/OPART
540 Gaither Road
Rockville, MD 20850

Name of Contractor: eShopperTools.com, Inc., DBA CleverSet, Inc.
Principal Investigator: Jane Jorgensen
(541) 829-6001 (cell); jorgenj@cleverset.com
Business Address: 673 NW Jackson Ave., Corvallis, OR 97330-4832
Phone Number: (541) 738-1010; (541) 738-6736 (fax)

Effective Date of Grant: 15 August 2004
Short Title of Work: FINAL REPORT

Contract Expiration Date: 31 Jan 2005
Reporting Period: 15 August 2004 – 31 Jan 2005

1. Specific Aims of the Phase I Grant

The research objective of the proposed work was to fuse in-home sensor data using relational Bayesian networks to infer behavior of occupants in a home setting, that is, to assess Activities of Daily Living (ADLs), and to detect events requiring intervention (e.g., falls). The sensor data were collected by CareWheels Corporation, a Section 501(c)(3) non-profit public benefit corporation. The CareWheels research project, "Internet-enabled Assistive Technologies for Independent Living and Aging-in-Place" has implemented a home sensor testbed currently operating at an independent living facility for people with severe physical disabilities in Portland Oregon.

The Phase I specific aims were to:

(1) develop a reasonable set of profiles that described suites of sensor outputs associated with types of behaviors. We anticipated that these profiles would reflect the temporality of behaviors (getting up in the morning, daytime activities, getting ready for bed), space utilization behaviors (living room behaviors, kitchen behaviors), or resident specific behaviors (wheelchair vs. walker behaviors). These models were be vetted by the CareWheels expert. This was intended primarily as a data-exploration / familiarization step;

(2) develop relational Bayesian network models that could infer behavior (ADLs and events requiring intervention) by modeling the joint probability distribution over the sensors, objects, and actions in the home; and

(3) assess the performance of the models using qualitative and quantitative analyses to evaluate the reasonableness of the models based on expert knowledge and the ability of the models to quantitatively infer behaviors from sensor data.

2. Materials and Methods

2.1 Data

The data were provided by CareWheels in the form of log files produced by CareWheels sensor system. The files were flat, delimited text files, with no relational structure. The files ranged from a week to four weeks of data with up to 130,000 sensor reports per file.

To convert the data into relational format, an *ad hoc* PERL program was used to parse the flat text files and reform them into xml elements appropriate to the schema being used for analysis.

2.2 Schema Construction

A series of increasingly comprehensive schema were constructed during Phase I. For the preliminary familiarization phase the schema identified *reports*, *sensors* and *rooms* as the relevant objects of discourse. Activities and events were not specifically named in the schema. Rather, they were modeled as probabilistic spatio-temporal clusters of reports.

Finally, as part of the Phase I activity we began development of a more complete ontology (set of terms and relations) that might serve as the basis for developing a complete schema for ADL and ERI detection and assessment. The ontology, developed in part from study of an earlier in-home monitoring ontology developed by Honeywell, the Consolidated Home Ontology in Protégé (CHOP) discussed in Section B.3.4 (Haigh et al, 2004).

2.3 Modeling Environment

Models were created and evaluated using CleverSet Modeler™. Modeler is a robust java-based research platform capable of constructing and evaluating dynamic and static RBNs. Modeler is capable of expressing arbitrary variables in relational data using a patent-pending synthetic variable language. Modeler is also capable of discovering RBNs from data.

2.4 Model Formulation

Several levels of models were constructed in the course of Phase I.

- **Naïve Bayes**. Naïve Bayes models are typically used in classification applications. They are Bayesian networks consisting of a single hidden variable representing an unobserved underlying state that captures correlations across sensor observations as "clusters," or values of the hidden variable.

- **HMM**. Naïve Bayes models, as formulated above, model the situation at a single point in time. Dynamicism was introduced into the models by extending the naïve Bayes model into a HMM with a single hidden variable. Hidden Markov models are dynamic probability models in which observables (sensor reports) are presumed to be explained, at least partially, by a "hidden" or unobserved state of the system (e.g., whether the client is moving about or stationary) which take into account only the previous state of the system when processing the current time frame.

- **DRBN**. Hidden Markov models represent the situation as a static, finite, and pre-enumerated set of possibilities whose actual occurrence may vary over time. This works well for small problems, but doesn't extend readily to complex problems where the number of possible situations is not easily enumerable in advance. Dynamic Relational Bayesian Networks (DRBNs) permit the modeling a situation as composed of any number of instances of defined *entities* (e.g., people, rooms, wakeup events) and *relations* (e.g, residence by a person in a room, movement between rooms). Our initial thesis was that this level of modeling would be needed to effectively detect and assess ADLs and Events Requiring Intervention (ERIs).

3. Results and Discussion

In this section we discuss results related to each of our three objectives (initial profiling of sensor data, inference of ADLs and ERIs, and assessment of model performance) in turn.

3.1 Profiles of Sensor Outputs Associated with Types of Behaviors

Profiles were first derived using naïve Bayes models. We began here because we hoped that a simple model would be expressive enough to discover some level detail about activity. We found these naïve Bayes models to be adequate at capturing broad swaths of activity which were not necessarily localized in time, but due to the structure of the sensor data were localized in space. For example, a hidden state representing inactivity in the apartment was correlated with both nighttime sleeping activity as well as absence.

We found these naïve Bayes models to be too crude to effectively infer ADL-level activities. They were unable to distinguish nighttime sleeping activity from when no one was home. Frequent activities tended to be represented by several hidden states precluding clear differentiation of activities.

3.2 Inferences about Behaviors

Since the data were inherently dynamic, it seemed sensible to use dynamic models instead of static naïve Bayes models. HMMs were chosen because they are easily interpretable and a standard model, the first line of dynamic model in a problem such as this. The advantages of the HMM are that they capture activity through time. Furthermore, by manually engineering the models, we can define specific activities as states of the hidden state variable. This, however, also proved to be a disadvantage: reasonable results were obtained only after extensive and time-consuming tuning of probability distributions of the hidden variables. One problem with the HMM approach is that it applies at the raw sensor report time-scale, a scale inappropriate for describing broader-scale activities such as ADLs. As a result, we were only able to effectively model the lowest level of activity.

DRBN models differ from HMMs in several important ways. First, target activities are represented not as mutually exclusive states of a hidden variable, but as separate data types. This means that we can consider several possible activities simultaneously. This is accomplished by triggering the generation of a new instance of an activity and subsequently determining its plausibility using the DRBN. Since the activities are a data type, they can have various attributes which can be incorporated into the Bayesian network. For example, wake-up activity may have attributes such location (e.g., from a nap in the living room), start time and end time. A second advantage of DRBNs over HMMs is that the former also allows for easier collection of statistics about activities. This is because each activity has its own associated set of detected instances stored in a relational database. Standard database queries can be used to evaluate, for example, number of hours of sleep per day.

Performance evaluation consisted of two parts: qualitative evaluation by the domain expert, and quantitative comparison of results against ground-truth and a competitive, commercially available system.

In the qualitative evaluation, the domain expert was asked to evaluate whether or not the results of the naïve Bayes, HMM and DRBN models produced were consistent with activities contained in the sensor data. In all cases, the DRBN results were deemed to be reasonable and reflective of the expert's perception of these situations as an active witness and as interpreted through the sensor logs.

3.3 Event Requiring Intervention: Detecting a Fall

Falls are rare events, and none occurred in our datasets. To evaluate DRBN ability in fall detection, our sensor expert manually modified a one-month data-set by inserting four simulated falls. The table below summarizes performance.

Hypotheses	**Detected**	**P > .1**	**Avg. Probability**
True Falls	4	4	0.9
Apt Exit	11	11	0.85
False Alarm	121	1	0.02
Total	136	16	

A total of 136 fall hypotheses were triggered over the one-month test period, an average of 4.5 hypotheses per day. Of these, the vast majority had posterior probabilities below 0.05. Of those with posterior probabilities above 0.1 (all had probabilities > 0.7), 4 were true falls and 1 was a false alarm. The false alarm corresponded to a period of time during which the bathroom was occupied continuously for over 1.5 hours, a highly unusual event. We believe this event could have been filtered out by more sophisticated models that, for example, fused activity information (the most likely activity throughout the period was "active") as well as time-of-day conditioned room occupancy expectations.

The remaining 11 events with probability greater than .1 were unexpected: our fall model correctly identified periods of non-occupancy. Since these are easily distinguished from true falls (the location is "front door outside" rather than an interior room), we view these as correct detections of non-occupancy rather than fall false alarms. Occupancy is a well-known "hard problem" in residence sensor monitoring. While we didn't evaluate the missed-detection probability for non-occupancy[†], we did verify that all detected periods of non-occupancy were detections of true events.

This demonstration of feasibility supports three core conclusions:

(1) Relational Bayesian networks (RBNs) are practical models of activity. RBNs consist of three central elements: data, schema, and probabilistic model. These three components support one another. The schema describes the relationships among data as they occur in real life. The probabilistic model is learned from the data, guided by the schema. RBNs are a concrete characterization of activities as they occur. In addition, RBNs support composition of higher-level ADLs (e.g., wakeup) from lower level activities (e.g., sleep, motion, etc.)

(2) RBNs provide logical models of activities of daily living. In general terms, RBNs aggregated data in ways that were consistent with general patterns of behavior with respect to time and place. On a more detailed level, using RBNs, we were able to track one particular activity, awakening, and one ERI, fall.

(3) RBNs provide an accurate interpretation of the sensor data in terms of activity. RBNs detected the wakeup activity accurately with respect to ground truth and their performance exceeded that of an alternative, commercially deployed system.

4. Summary of Progress

The Phase I work accomplished the Specific Aims of the proposed work and achieved project milestones on time.

[†] Evaluation of the missed detection, or false-negative, rate would have required manually reviewing the entire one-month sensor data-set yet again to identify all true periods of non-occupancy.