

The BehaviorScope framework for enabling ambient assisted living

Athanasios Bamis · Dimitrios Lymberopoulos ·
Thiago Teixeira · Andreas Savvides

Received: 31 December 2008 / Accepted: 5 October 2009 / Published online: 3 March 2010
© Springer-Verlag London Limited 2010

Abstract The in-house monitoring of elders using intelligent sensors is a very desirable service that has the potential of increasing autonomy and independence while minimizing the risks of living alone. Because of this services that improve safety and quality of life is a very promise, the efforts of building such systems have been attractive proposition for elders living alone. Although the spanning for decades, but there is still a lot of room for improvement. Driven by the recent technology advances begun to become more relevant for two main reasons. First, in many of the required components, in this article, we present a scalable framework for detailed behavior interpretation. Our framework supports in-house monitoring of elders using an intelligent gateway and a set of cheaply available sensors, in addition to more advanced camera-based human localization sensors and client for GPS-enabled mobile phones that provides monitoring when outdoors. In this article, we report our experiences and present our current progress in three main components: sensors, middleware, and behavior interpretation mechanisms spanning from simple programmable rule-based alerts to algorithms for extracting the temporal routines of individuals.

Keywords Assisted living framework · Activity recognition · Activity learning · Assisted living interfaces

A. Bamis (✉) · T. Teixeira · A. Savvides
ENALAB, Yale University, 51 Prospect St. Room 000,
New Haven, CT 06520, USA
e-mail: athanasios.bamis@yale.edu

D. Lymberopoulos
Networked Embedded Computing Lab,
Microsoft Research, Redmond, USA
e-mail: dlymper@microsoft.com

Monitoring people activities and providing automated services that improve safety and quality of life is a very attractive proposition for elders living alone. Although the problem was considered for many years, it has recently begun to become more relevant for two main reasons. First, many studies together with the rising costs of health-care present a scalable framework for detailed behavior interpretation. Our framework supports in-house monitoring of elders using an intelligent gateway and a set of cheaply available sensors, in addition to more advanced camera-based human localization sensors and client for GPS-enabled mobile phones that provides monitoring when outdoors. In this article, we report our experiences and present our current progress in three main components: sensors, middleware, and behavior interpretation mechanisms spanning from simple programmable rule-based alerts to algorithms for extracting the temporal routines of individuals. From a technology perspective, the majority of components required to build such systems are becoming readily available. Many systems under development both in academia [7, 14, 16, 20] and industry [2], as well as some commercial systems [4, 8] are already capable to provide essential monitoring services (for a survey of current state-of-the-art see references [3, 6]). What is mostly missing is experience and systematic knowledge to intelligently assemble the components into robust architectures and practical, deployable systems. In addition, most of these systems focus on collecting and presenting simple statistics, often using intrusive sensors (e.g., wearable devices), requiring, thus, the involvement of health-care providers and stakeholders in the system loop.

The BehaviorScope project at Yale [9] is investigating these challenges by trying to build a functional system that can autonomously understand behaviors with enough detail to provide meaningful services. The goal of the project is to design an extensible architecture that can use a wide variety of sensors to interpret human activity, dynamically

generate activity models, and use them to generate alarms, reports, triggers, and to answer queries. In this article, we provide an overview of the architecture of the system under development, and report on the main components that our research is trying to address. To this end, Sect. 2 provides an overview of our system requirements. Section 3 outlines our system architecture and Sect. 4 explores various methods, we have considered for interpreting the data. Finally, Sect. 5 concludes the article.

2 Overview

The provision of services requires a set of sensors to be deployed inside a home to observe the inhabitants, interpret the observations, and provide meaningful responses. Depending on their condition, one can anticipate that the home inhabitants would be willing to subject themselves to a certain level of observation (i.e., give up some of their privacy) in exchange for services. The goal of our architecture is to provide a versatile system that can accommodate this at different levels, from very simple to very detailed observation, according to individual needs. The initial form of the system is intended for elders that live alone, and are fairly independent. In this case, the role of the system would be to eliminate certain risk factors that could otherwise be avoided by resorting to institutionalization. In its simplest form, such a system would offer a wide variety of services:

- *Queries* The system should be able to answer queries such as: where is the person, is that person getting enough sleep, is the person out of the house beyond the expected time?
- *Alarms/notifications and triggers* Notify stakeholder when the person returns/leaves the house, notify when the person wakes up/goes to bed, generate an alarm when the person did not wake up within a threshold after its usual time.
- *Detect anomalies* By observing and learning routines (e.g., daily, weekly, monthly), the system can provide notifications when an unusual deviation from the routine happens.
- *Recognize specific behaviors* By allowing the programming of specific behavior recognition libraries into the system, one can tailor the system to provide customized observations and actions for each house. This, for example would help tailor the same system to people suffering with cognitive decline and people who are frail and run the risk of falling or getting stuck somewhere (e.g., bed, toilet).
- *Actuate* Take action when certain events (or combinations of events) are detected.

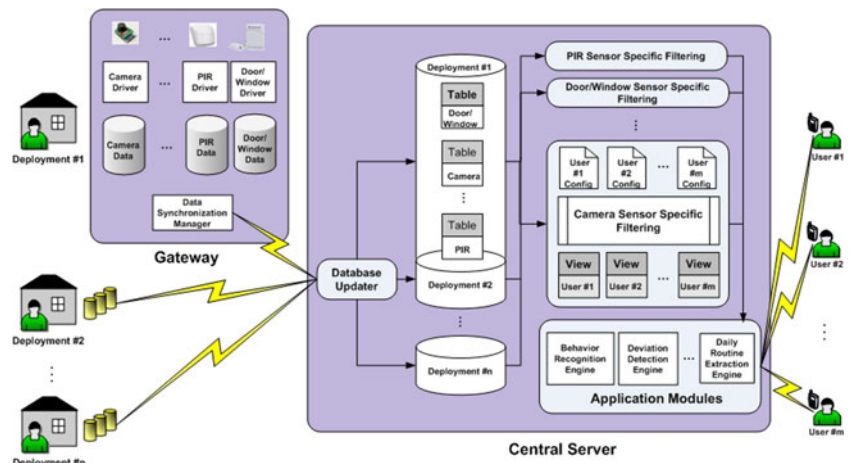
The users of the system should be able to configure the above properties to adapt it to their individual needs by programming custom triggers, defining custom queries for future use and specifying what actions should be taken when a specific behavior is detected. Moreover, for detecting routine behaviors and timing parameters, the system should be able to use a generic specification as a starting point and automatically “cast” itself to the home and the activity patterns of individuals when it is actually deployed.

The above requirements create a new set of challenges involving sensing and data interpretation, and call for a middleware architecture that can support a heterogeneous set of devices and their tailored configuration for each home. Furthermore, for cost effectiveness and ease of installation, a practical system should provide the aforementioned services without requiring the exhaustive tagging of every item in the home with sensors. To make this possible, the BehaviorScope project seeks to build a rigorous understanding of what today’s off-the-shelf sensors can do, what types of new sensors are required and how a heterogeneous set of such sensors can co-exist in the same framework to collect and interpret data.

An outline of our system-wide architecture is shown in Fig. 1. A set of wireless sensors is placed at key locations to collect sufficient information for recognizing a person’s activity profile around the house. The data collected by the sensors is forwarded to an intelligent gateway installed inside the house that processes and interprets the data by communicating with a central server. Currently, our system supports a wide range of sensors spanning from widely available motion sensors to privacy preserving camera-based human localization sensors that we have been developing as part of the BehaviorScope project. In addition, our system provides a client for GPS-enabled mobile phones (currently supporting BlackBerry smart-phones) that can track elders while outside the house and provide additional services, as a virtual escort service and location based reminders.

Caregivers and stakeholders can interact with the system via two main interfaces, a mobile phone interface and a web interface. The mobile phone is the main interface for communicating daily summaries, alarms/notifications, triggers, and queries. The web interface supports a more elaborate setup that allows the end-user to customize the behavior of the system to each home. In particular, the web interface allows to view the incoming information (and historical information depending on the privacy requirements) using text and graphically using plots (both simple intuitive plots and more advanced, information-rich plots), share the information with other users, setup very simple alerts/notifications through a graphical interface, and program their own complex triggers using STFL, a close to

Fig. 1 General overview of the system architecture. By defining a modular architecture, and multiple levels of abstraction, we can achieve scalability and robustness




human scripting language. Moreover, the system can provide simple statistics as the “motion level” of the person according to the specific monitoring needs. Although it (also referred to as circadian rhythm), extract simple statistics or run more complicated algorithms in order to extract the spatiotemporal routine of the person and detect anomalies. Finally, given a proper definition, a hierarchy of Probabilistic Context Free Grammars (PCFG) can be designed for detecting more complex human behaviors, as

3 System architecture

A key premise of the BehaviorScope infrastructure is the ability to jointly consider information from multiple sensor types to infer behavior from low-level data. Most of the sensors are off-the-shelf passive infrared (PIR), and door/window sensors, assisted by more powerful, motion detecting and discriminative sensors derived from cameras. The latter forms of sensors is aiming to define a new sensing modality in which people locations and movements in the house can be sensed but no images can be produced. The home is divided into two types of areas, common and private (see Fig. 2). Counting sensors are only placed in the common areas of the home especially near the exits. PIR sensors and

Fig. 2 The areas of a house are separated into common and private. Common areas can be monitored by “intrusive” sensors as cameras, whereas private areas can contain only “non-intrusive” sensors as PIR



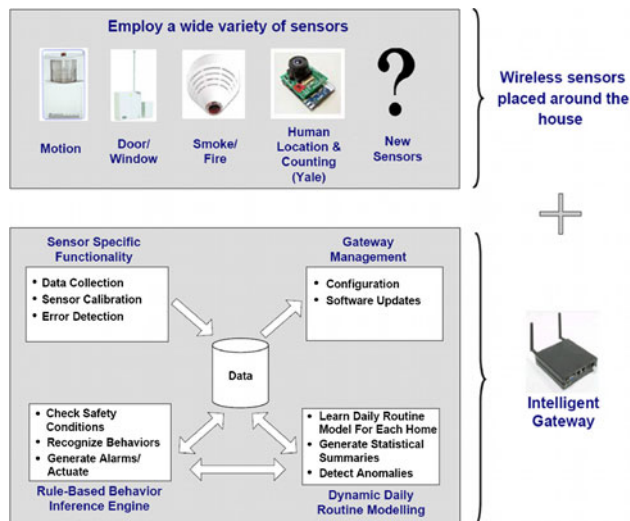


Fig. 3 A home sensor network kit consists of one or more sensor networks and an intelligent gateway that can manage the sensor networks, collect data, pre-process them and transmit them to the Central Server

dump it into the gateway's database. The system will automatically take care of the data synchronization process with the central server. In addition, the developer can add a number of preprocessing modules, depending on the type of "fundamental" information that needs to be extracted from the data. For example, in the case of a camera node, we can define areas of interest and generate an event whenever motion is detected inside that given area. The outputs of the preprocessing modules are added back to the database, and can be used by applications running either on the central server or locally by the end-users. Similarly, in order to add a new application on the central server, all we need to do is add a preprocessing module, that given possibly some configuration parameters from the users, will convert raw (or previously preprocessed) data into the proper format and it will then pass them to the application module.

3.1 Camera-based privacy preserving counting and human localization sensors

The PIR and door sensors used in our system are off-the-shelf sensors readily available from different vendors. Although PIR sensors detect motion, they do not necessarily detect occupancy of an area inside the house. For instance, if two people enter a room, and one person leaves, our measurements from PIR sensors alone cannot easily determine that a person is still in the room if the person does not move. Moreover, most commercial sensors have very primitive MAC layers, primarily geared towards the security alarm trigger applications. This does not always favor assisted living setups where readings from multiple

sensors, and their relative timing have a meaning. A straightforward solution is to attach PIR sensors to off-the-shelf sensor nodes, but that would cancel their main advantages of low cost and increased battery lifetimes.

Because of these limitations of PIR sensors, and the need to count and track multiple people, we are currently developing a new custom sensing modality that can localize and track people inside the house without requiring them to wear a tracking device. Although the sensor is derived from cameras, it directly aims at the development of a new camera chip that can localize, count and track people without providing any image information to the rest of the system.

Due to the extensive amount of processing that is generally required for computer vision tasks, camera nodes architectures in the literature have typically followed one of two approaches: (1) using the fastest low-power CPU available; (2) adding specialized processing components that are capable of a high degree of parallelism, such as DSPs or CPLDs. In our research, we take a third route, by making fundamental changes to the underlying image sensors themselves. The typical image sensor outputs a serialized array of pixel intensity values. This array contains raw data that must be heavily processed before any desirable information can be gathered. What is more, only after the image is processed can one know whether or not the captured scene is interesting. The result is that many uninteresting frames end up being captured and processed before ultimately being discarded, resulting in a large waste of resources.

Our platform, on other hand, is built with biologically inspired address-event (AE) imagers in mind. Instead of outputting arrays of pixel intensity, these imagers asynchronously output an *address* (in pixel coordinates) every time an *event* is detected. Events can be any measurable phenomenon. In the case of the imagers we use, an event is triggered every time a pixel senses motion (an above-threshold change in intensity). The power of address-event lies in three separate properties: processing occurs at the pixel level, freeing the controlling CPU from complex imaging tasks; AE sensors do not discretize time into "frames", which allows for precise measurements and provides privacy; AE sensors are typically ultra-low-power.

In our current platform, we emulate the address-event imager in software. The emulated parameters are used to guide our custom hardware imager design. Since our algorithms are written for address-event input, once a hardware AE design is fabricated it can directly substitute for the emulated version. The sensor nodes in our deployment use Intel iMote2 sensor nodes coupled with a custom camera board. The purpose of the nodes is to find and track the people in their field-of-view, communicating the detected coordinates back to their base. The software on the sensor nodes detects humans based on size and motion by constructing a motion histogram

[18]. The histogram utilizes person-sized bins to compute a density estimation of possible human locations. This is done by dividing the image into partially overlapping person-sized areas, and counting the number of threshold motion pixels that lie within each area. These counts are organized as bins in a two-dimensional histogram, and the local maxima are computed to locate the histogram peaks. Each peak indicates the likely location of a moving person, as seen in Fig. 3.

3.2 Intelligent gateway

Our gateway architecture consists of four main categories of software modules, shown in Fig. 4. The first type of modules are sensor specific and consist of the drivers receiving the sensed data from the network, removing correcting erroneous measurements, detecting malfunction of the sensors and, generally, managing the correct operation of the deployment. The collected data is stored in a local database, which is incrementally (i.e., only new data) transmitted to the central server, by a synchronization module. Other modules in this category include modules for receiving software updates, modules for checking the correct functionality of the gateway, modules performing authentication, modules allowing the remote configuration of the gateway parameters and, generally, any module that is not sensor specific or concerns data processing. When it comes to data processing, there are two categories of software modules. The first category involves modules that collect statistics, learn from the collected data and possibly respond on significant deviations, whereas the second category includes modules that try to detect certain behaviors and patterns inside the network and possibly take certain actions as a response.

3.3 Central Server


Besides the system management modules (e.g., gateway software updating module) and the module that updates the database with the incoming data from the deployments, the central server contains two more main categories of software. The first one has to do with preprocessing and conditioning of the incoming data, and depends on the type of the sensors and the requirements of the end users. In the case of complex sensors, such as cameras, the data collected from every type of sensor can be processed in order to extract some features that can directly be used by the users, or be given as input to one or more applications. The central server stores the data in a separate database for each deployment and incrementally preprocesses the data according to sensor types and the required information that needs to be extracted (i.e., according to the data processing module that we wish to use). The new data is then passed to a sensor specific module, which using user-specified and statistically learned configuration parameters creates views for each user and each data processing module that is available for the given deployment and sensor. These views are subsequently accessed by the application modules located inside the server, which will generate a number of results, or by custom applications designed by the users (and located outside of the server).

3.4 The BehaviorScope web portal

To be of use, every assisted living environment must provide both a synchronous interface for real-time monitoring of the persons of interest and an asynchronous interface that can be used to communicate notifications, updates and most importantly alerts in cases of emergency. Since the interface needs to be accessible to a large and diverse set of users the best option is to implement it as a web service, where people can login securely, access all the information they need and configure the types of notifications, updates and alerts that they require.

This set of users can include, besides the monitored person, one or more stakeholders, caregivers, emergency personnel, persons that reside close to the monitored person and can be of assistance in cases of emergency), the researchers, social workers and others. Each of these groups has in general different requirements from the

Fig. 4 Multiple people counting and tracking can be achieved by estimating a motion histogram and detecting its peaks



assisted living environment, and needs to access different provides different forms of visualization of the data (see for types of information. Moreover, the monitored person will example Fig.5), different types of statistics and advanced usually be willing to sacrifice different amounts of privacy interfaces for the customization of a deployment, as well as in exchange for services provided by each group. To what the capability to define simple email and SMS based is more, information meant for different groups of users notifications and alarms (see Fig.6).

will usually involve different levels of anonymization. Every datasource is characterized by a serial number, Hence, the interface for an assisted living deployment which identifies it uniquely and allows web-based record-should provide many different types of representation of and access to the data. Every deployed sensor network the collected data to accommodate the needs and technology appears as a different datasource on the portal and a competency of different users. Furthermore, it should gateway can support transparently different types of network-provide different types of statistics, and it should allow works using different serial numbers. In addition, a user is programming different types of notifications, alerts and allowed to define a new datasource using some primitive statistical summaries. These services should be provided by data types offered by the system and re-use the existing both a graphical environment and using natural language data sharing, visualization, statistics extraction, and data In addition, the interface should provide both online access interpretation mechanisms.

to the incoming data and access to the history of the particular monitored person, depending on the privacy constraints. Configuration and customization options are communicating alerts, high-level statistics and enabling the essential to allow the users to share information mainly user to perform simple queries.

taining full control of the type and amount of information that different users can access.


The BehaviorScope web portal (BScopeWeb) provides most of this functionality by allowing users to register

datasources of different types (currently: motion sensors, localization sensors based on cameras, rfid, GPS-enabled phones, and simple sensors as door/window, tamper or temperature/humidity sensors), and share these datasources with other users of the system. BScopeWeb

3.5 A SpatioTemporal Itering language (STFL)

for enabling actuation

Fig. 5 Occupancy information for the rooms of an assisted living deployment. This visualization type can be used for easily inspecting the durations spend in each room for a given time interval



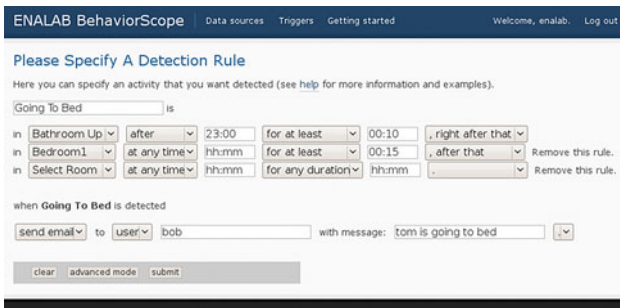


Fig. 6 The advanced user interface provided by STFL (STFL GUI) can be used by the users of the BehaviorScope system to define their own rules and sequences of rules and constraints to actuate upon the detection of events of interest

notifications and alarms to the users of the system are email messages and SMS. Other options include phone calls, updating visual notifications (e.g., turning on lights, changing the displayed image on a digital picture frame) or triggering audio alarms. Recipients of the alarms and the notifications can include the monitored person itself (e.g., “remember to take your medication”), stakeholders, caregivers, and emergency personnel. Besides these simple forms of actuation, the user is also allowed to programmatically create new types of “virtual datasources” that are generated from higher-level semantics extracted from the data. These datasources, among others, enable fine-grained control to the access rights of different users and groups of users to the data of the monitored person. All these become possible, by allowing the owner of a datasource to define simple rules and sequences of rules and constraints using the SpatioTemporal Filtering Language (STFL). [

STFL is a close to natural language that consists of four different layers (see Fig. 7), each aiming at users of

different expertise. The main assumption of STFL is that high-level activities can usually be decomposed into simple rules or simple Finite State Machines (FSMs) describing the activity in terms of sequences of locations and specific temporal characteristics. The two top layers can be used from within the BehaviorScope web portal, whereas the two bottom layers aim mostly at the developers of the system. In particular, the “Advanced User Interface” (STFL) provides a programming language close to natural language, which can be used to describe rules and sequences of rules and constraints, whereas the “Graphical User Interface” (STFL GUI) provides a graphical interface (see Fig. 6), through which users can set simple triggers. To this end, one of the most important aspects of STFL is that it can provide a large set of conditions for when an alarm will be triggered. For example, a user can through a simple interface specify the number of times that it can be triggered or time intervals that the alarm has to be activated or deactivated.

3.6 Beyond in-house monitoring using GPS-enabled phones

In many cases, extending the monitoring of a person outside of the house can both increase safety and provide valuable information about the condition and the routine of the person. The BehaviorScope system allows users to record their position using GPS-enabled phones. [Our initial deployment in an urban setting heavily relies on the client application running on mobile phones. The deployed mobile client (currently supports GPS-enabled Blackberry phones) can be downloaded from the BehaviorScope web portal [5] and supports several features that enable it to be an active contributor to the overall system architecture.

Since the goal of this application is to make it easy to stay safe and secure anywhere, at anytime, the application needs to efficiently manage its power consumption and make its state known to the server at all times. The application informs the server of its status on power up and shutting down, loss of GPS signal, and feature usage. To conserve power, local processing, intelligent sampling, and other sensors such as accelerometers need to be exploited. Our prototype experiences have shown that reading the GPS alone can take a noticeable toll on the phones battery lifetime. Such excessive power consumption could be reduced by utilizing accelerometer sensors and context inferred from the behavior monitoring applications to intelligently manage the GPS sampling and communication frequency. The BlackBerry smart phones used in our prototype deployment do not have accelerometers but other phones such as the Nokia N95 and iPhone already have them. We anticipate that more phone models will have them in the future. In addition, we are currently looking

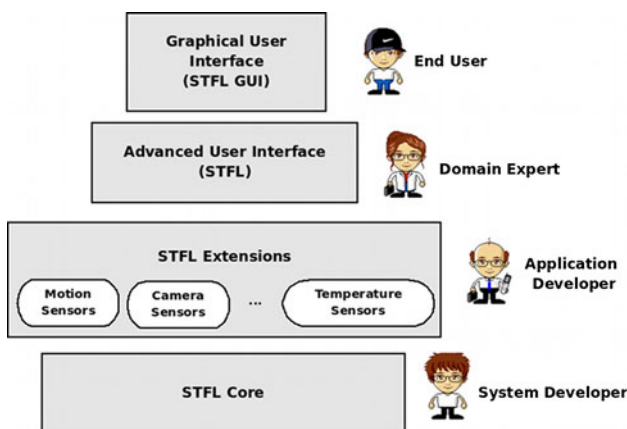


Fig. 7 STFL provides four different layers that aim to users of different expertise. The two top layers can be used by the end-users of an assisted living environment to define their own custom triggers, as well as by the owners of an assisted living deployment in order to control the access rights for other users and group of users

into the possibility of using signal strength variations or the availability of other forms of information (e.g., existence of wi networks) to determine if a person is indoors or outdoors and accordingly enable or disable the service.

The mobile client contains a basic tracking feature that users can select to turn on to allow their location to be sent to a central server. The user can choose to have this feature on 24-h a day, or just during select commutes or times of the day. Ideally, the application will be running continuously so as to collect as much information and allow as much personalization as possible. This location information is then accessible securely from any computer or mobile device that has access to the Internet via the BehaviorScope web portal.

A “Virtual Escort” feature is an integral part of the mobile client. It allows users to have an escort when the user cannot find anyone else to commute with. This feature provides a cost-effective and time-efficient solution to staying safe and gives the user access to a programmable PANIC button that can let interested parties know there is trouble and exactly where the user is. The web interface also allows users to set up triggers that inform family and friends via SMS or e-mail alert when the user is leaving or entering a pre-defined space. This automates the process of checking a user’s location by having an automated message sent out according to the pre-specified preferences. These triggers can be simple conditions about geographic locations (see Fig. 8) or more advanced rules and sequences of rules and constraints described using STFL.

If the user defines certain areas as being associated with specific activities, the system engine can write automatic digests of a user’s day to send to friends and families. This would allow a user to automatically “keep-in-touch” with the monitored person, even when the persons are very busy and have no time to call or e-mail.




Fig. 8 The BehaviorScope web portal allows users to specify locations of interest on a map and generate a message whenever the monitored person, carrying a GPS-enabled phone running the mobile client, enters or leaves these areas

4 Data processing and interpretation

Depending on the granularity of the data and the application of interest to the end-user, the system can provide a wide range of statistical information. In the following sections, we will use data collected from online deployments in two different homes, consisting mainly from PIR and door/windows sensors. *Deployment A* has been conducted in the USA for more than 7 months, whereas *Deployment B* monitors an elder couple and their adult son in Cyprus for the past 4 months. In both deployments camera sensors are located near the exits of the house and are used only for counting the number of persons present in it.

The following sections first discuss the statistics we can extract from motion-only information generated from PIR sensors in the BehaviorScope deployments. The discussion is separated into two cases, the case where we have a single person living in the house and the case where we have more than one persons living in the house. Afterwards, we are going to shortly discuss a method for detecting significant deviations from the “normal” living pattern of a person or a house, and in the following two sections, we are going to discuss how our system can automatically generate a high level model of the daily living pattern of a person, as well as show how it can be programmed to detect specific behaviors.

4.1 Motion statistics


The lowest level of information, we can extract from a motion sensor is a time-stamped notification of when motion was detected. Although PIR motion measurements are not sufficient to determine occupancy (i.e., whether a person is in a certain room or not), they can provide information about people movement inside the house. This information provides an indication of the room occupancy patterns inside the house.

4.1.1 Single person case

In the case where only a single person is in the house, time-stamped motion sensor measurements capture the room-to-room transitions of the person. This information on its own can reveal the activity profile of a person and the level of periodicity of a person’s daily routine. Figure 9 shows the room transition profile of the elder in *Deployment A* over the period of 1 week (September 19–25, 2007). The sequences reveal that the person has a very consistent daily pattern, and with a few basic rules and statistics we can extract basic activities and sleep patterns.

Moreover, using simple rules and collected statistics (e.g., average sleep duration) we can detect very simple activities, as for example night sleep. Night sleep in the

Fig. 9 Room transitions and detection of sleeping activity and bathroom usage (using rule-based triggers) of the monitored person of Deployment A (person living alone) for a period of 1 week



deployment of Fig. 9 can be inferred when we detect time in the living room between 3:30 and 6 p.m., some motion in the third bedroom of the house (BR3), after which cannot happen during a normal weekday when 11 p.m., and followed by lack of motion for at least the person is at work approximately 30 min.

Assuming that a person is not moving between two consecutive motion notifications, we can, additionally, provide occupancy statistics for a given location. In the case of multiple people living in a house, motion Figure 10a shows the occupancy of the living room of sensor data loses its sequence properties and cannot reveal *Deployment A* (person living alone) for the duration of a the daily patterns of one person in specific. Since, the week (August 9–15, 2007). From this figure it is easy to sequence property is essential for inferring occupancy observe that the person spends significant amounts of time in the living room, and usually around the same time of the occupancy pattern for the living room of *Deployment B* day. In particular, we can see that the person will always spend time in the living room (watching TV), late in the provides some useful information on the usage profile of morning and noon, until she goes to work around 3:30 p.m. To provide meaningful statistics comparing the utilization of the rooms of the house, we first need to define a common representation of the “quantity of motion” for a given area and time window, that we are interested. Hence, it is easy to observe, that during the weekend this pattern of the rooms of the house, we first need to define a common representation of the “quantity of motion” for a given area and time window, that we are interested. Hence,




Fig. 10 Occupancy of the living room inferred from motion information for *Deployment A* (single person living alone) and *Deployment B* (multiple people)

we define a new metric for the mobility of a person at a given location and time window, called *relative mobility level*. The relative mobility level is essentially the normalized amount of motion in a specific area and time window with respect to a given time period and a given area of interest.

More formally, given the minimum time duration t_{\min} (e.g., 15 min) for which we are interested we can define a set of *time bins* t_{b_i} over a given time period T (e.g., a day) as $t_{b_i} = [\frac{T}{i \cdot t_{\min}}, \frac{T}{(i+1) \cdot t_{\min}})$. Similarly, for a given set of sensors S we can define a set of *space bins* s_j as the union of space covered by one or more members of S (e.g., the area covered by sensors with ids 5 and 11, which is the living room area for Deployment A). If we take all the pairs of space and time bins, we define as *relative mobility level* of each such pair the total number of motion notifications that we received in the particular space bin and the particular time bin over the total number of motion notifications we received for all the period T and all the sensors in S .

For example, Fig. 11a shows the average relative mobility level of an elder person living alone in 15-min intervals. In this case it is defined to be the entire house (i.e., all sensors) and the time bins are selected to be 15-min durations during the course of the day. Consequently, every bar in the graph indicates the average “quantity of motion” at the given 15-min window during the day for the entire house of Deployment B, which is an indication of their utilization. We separate weekdays from weekends, in order end-users or specific applications require, we can have different types of resolution. For instance, in Fig. 11b

From Fig. 11a and b it is easy to extract useful information for the daily living pattern of the monitored person. From the plots it is easy to infer when the person is sleeping or is out of the house by combining measurements with other context information such as the time of the day




Fig. 11 Relative mobility levels for the entire house and Deployment A for time bins equal to 15 min, b Deployment A for time bins equal to 1 h, c Deployment B for time bins equal to 1 h




Fig. 12 Average daily relative mobility level for every room of Deployment B. Weekdays are separated from weekend days in order to demonstrate that they follow different occupancy patterns

same. Similarly, we can see that during the weekend the family spends significantly more time hanging out in the *Sun Room* (“SR”) of the house and, also, more time in the *Dining Room* (“DR”) having lunch and dinner.

4.2 Detecting deviations

To detect deviations in the living pattern of a person or the house, we have first to model the motion activity pattern in a way that will enable us to find regularities, thus defining a notion of “normal”. Apparently, every person or house (in the case of multiple people) has its own pattern, which changes over time and is also dependent on many macroscopic parameters, as for example the time of the year or in the case of the house, the current set of people living inside it. Of course, we can try to detect deviations in many different time windows, but for the following discussion we will limit ourselves to detecting “deviating days”. Besides a daily pattern, a person or a house can have patterns in many different time resolutions both larger and smaller. For example, most people have a certain wake-up routine and a house has a yearly usage pattern, that is, the utilization of the rooms changes depending on the season of the year. This becomes apparent even by simple observation of Figs 10, and 12, where it can be seen that the normal pattern of a weekday presents several differences from the pattern of a weekend day, both when we focus on the pattern of a person and when we focus on the pattern of a house.

To learn the daily pattern of a person, ideally, we need to discover what remains invariant, possibly adapting over time to the new parameters. Our intuition is that a person will spend approximately similar amounts of time in a place over the course of a “normal” day and will produce

proportionally equal amounts of motion information. By modeling the daily motion pattern of a person using a vector, with each field of the vector indicating the relative mobility level at a given place of the house during a given time window, we can use the distance of the vectors as an indication of how different 2 days are. We expect that the distance between “normal” days will be relatively small in comparison to “deviating” days. Thus, we can define as “deviating days” any vectors who are outliers. In order to detect outliers a clustering algorithm, such as k-means, can be used.

4.3 Extracting activity models

Moving to a different category of information, instead of just trying to collect statistical information or identify variations of the daily living pattern of a person, we can try to detect repeating patterns, and based on them create a model of the daily habits of the person. To accomplish this, first, we model the sensor network as a spatiotemporal symbol generator that is triggered by the monitored person as she moves over space and time. Based on our network model, we formulate the problem of finding the daily activity model of a person as the problem of finding the most probable, network-level, sequences of node-level, sensing features, namely location, time and duration. By the simple observation of Figure 8 it is, already, easy to observe that that the daily activity of the person under observation has regular recurring patterns.

While the statistical representation of the raw sensing data and its variation over time can provide valuable information about the monitored person, it fails to provide an in-depth analysis about the person’s daily living habits. As the person moves inside the house, a sequence of detected sensing features is produced over time. These features might encode spatial information, such as the rooms/areas the person visits or the objects with which she interacts, as well as temporal information, such as the exact time and duration of these features. The sequence of these recorded sensing features over the course of a day represent the monitored person’s daily activity signature. Using this stream of symbols, we formulate the problem of human activity modeling as a spatiotemporal pattern-matching problem on top of the sequence of recorded sensing features and solve it using an exhaustive search algorithm [10]. The exhaustive search algorithm we use is based on the a-priori principle: any subsequence of a frequent sequence has to also be frequent. Given this, we have employed an

exhaustive, yet very efficient, algorithm that automatically discovers the most frequent sequences of sensing features. Initially all the frequent sequences of size 1 are discovered. Then, using the set of frequent sequences of size 1 as our starting point we identify the most frequent sequences of


size 2 and the algorithm continues iteratively until no activity is associated with two temporal parameters; the frequent pattern is finally discovered.

At the end, the most frequent sequences of features activity and its duration (e.g., sleep lasted 8 h). different sizes have been identified. Since these sequences represent the monitored person's frequent activities, characteristics lies on the fact that time and duration of a sensed combined, they can be used to build the daily living model event, are continuous variables that can take any value. of the monitored person. For instance, Fig. 13 shows the To consider them in a model, these quantities need to be daily living activity model that was extracted out of appropriately discretized. In doing so however, one needs 30 days of recorded data of an elder person living alone into consider the fact that temporal characteristics may Deployment A. In this case, the basic sensing features differ in two ways. Temporal variations within one par-recorded were very primitive activities such as sleeping, particular event's time, and temporal variations across different having breakfast etc.

4.4 Learning temporal characteristics

Although reasoning with sequences of events and locations can reveal many informations about a persons routine, it lacks a significant component characterizing every human activity, namely time. The stream of symbols generated by the sensor network contain a temporal dimension, which can be used to improve our knowledge and increase the accuracy of our models. In particular, every event and source is needed. The goal of this process is to provide an

Fig. 13 Simplified daily living model of the monitored person of Deployment A



answer to the question: “when and for how long does this 175 min and one, significantly less probable, starting at event type take place?” about 16:35 lasting for about 190 min.

The BehaviorScope framework, provides a method for extracting the temporal characteristics of an event or activity by formulating the problem as a clustering problem

of two-dimensional vectors and proposing a new metric and a new agglomerative clustering algorithm for grouping together events and activities with very similar temporal characteristics [10]. Every vector contains the start time and the duration of an event, which essentially correspond to a line in the plane (see Fig. 4). A distance metric called pairwise-density is used to characterize the degree of overlap between two lines or equivalently how similar the temporal properties described by two vectors are. Finally, an agglomerative clustering algorithm attempts to maximize the overall average pairwise-density of the clusters, without making any assumptions about the number of clusters.

The output of this algorithm (see Fig. 4) are clusters of the instances of an event or an activity with very similar temporal characteristics, along with a metric indicating confidence for the particular cluster. Using the latter metric, we can identify relatively invariant activities of the person (i.e., activities that will repeat at approximately the same time and for the same duration every day) and improve our model by defining new symbols that incorporate time. As an example, Fig. 4 shows when the monitored person of Deployment A leaves the house. From this figure, we can identify two main clusters, the one starting at approximately 15:15 and lasting for about




Fig. 14 The different instances of the “Out” activity (i.e., the person being outside of the house) of the person of Deployment A for the period of approximately 1 month. With the same color appear clusters with instances with approximately the same temporal characteristics (i.e., start time and duration). In the plot the x-axis depicts different dates and the y-axis depicts absolute time for a given date. Thus the start of a line indicates the start of an instance of the “Out” event and the end of the line symbolizes the end of the event

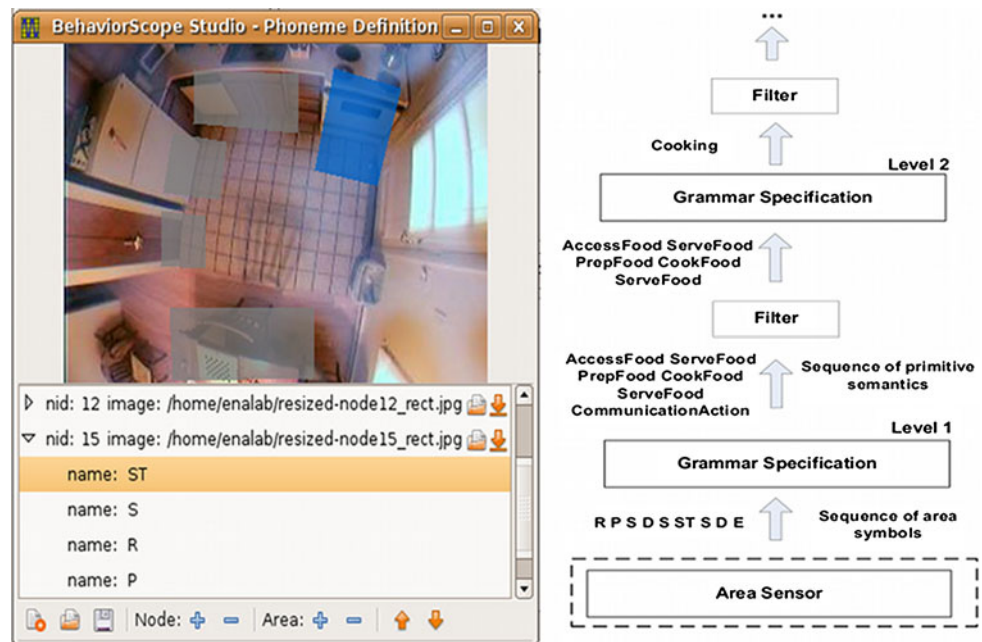
4.5 Rule-based activity inference

our project has also developed a behavior interpretation system with which users are able to describe activities as a collection of probabilistic rules with spatial and temporal characteristics expressed in high level script form. Each activity description has well-defined inputs and outputs enabling the creation of a library of activity components that can be connected together into hierarchies to provide even more complex interpretations. The power of such a framework comes from the hierarchical organization of reasoning. This allows the use of simple timestamped, localized sensor measurements to reason about more macroscopic behaviors taking place in space and time.

The main idea is that human behaviors are sequences of very primitive actions that take place over space and time. Different activities can be described by simply combining these primitive actions over time in different ways. A multi-modal wireless sensor network monitoring a person’s location and interaction with different objects over space and time provides a stream of basic sensing features for identifying these primitive human actions. The proposed method suggests to parse the sequence of detected sensing features into higher level human behaviors in a hierarchical bottom-up processing model that is similar to natural language processing. The set of recorded features becomes the human activity alphabet. In the same sense that we combine letters to form words, we combine these features to define primitive actions; similarly, as words are combined to form sentences, sequences of primitive actions are combined to describe basic human activities; and so on from sentences to paragraphs, paragraphs to stories, we combine human activities over space and time to define macro-scale human behaviors.

The basic interpretation blocks in this hierarchy are probabilistic context-free grammars (PCFGs), [12] that can be either specified by the user or even automatically extracted from the collected data as shown in [10]. Through a simple high-level interface, users provide a collection of probabilistic rules that form a PCFG. This set of rules specifies one or more activities by enforcing a syntax on the recorded input stream of sensing features (for an example see Fig. 5). This syntax takes into account spatial characteristics (detected sensing features and their sequences over time) as well as temporal characteristics. A flexible time abstraction layer we have designed and implemented [13], enables users to associate time information to the recorded sensing features on a per-grammar and on a per-feature basis allowing the definition of

Fig. 15 To detect an activity the user must a) specify a model of the areas and events of interest, and then b) specify a hierarchy of Probabilistic Context Free Grammars (PCFG) that specify the activity that needs to be detected



grammar specific spatiotemporal features. By parsing these wear sensors. Furthermore, we have discovered that there is a lack of synergy between learned and predefined models. Our work up-to-date has demonstrated that the two model types are complementary, and in order to deploy an effective system the two models should work together in close coordination. This and the development of intelligent motion discriminative sensors will become the focus of our future work.

The grammar hierarchy interpretation framework has already been used in several home network deployments to automatically interpret the recorded stream of data and provide meaningful activity summaries [1, 13]. Its interpretation power has also been demonstrated by the successful detection of complex activities, such as the cooking activity [12].

Acknowledgments This work was partially funded by the National Science Foundation under projects CNS 0626802 and CNS 0751513. Any opinions, findings and conclusions or recommendation expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).

5 Conclusions and future work

This paper described our up-to-date progress on a large scalable system for monitoring elder activities in assisted living. In particular we presented, the current status of the various components of the system, including among others our work on the development of new camera-based sensing modalities, the development of middleware for collecting, transmitting, sharing and visualizing data from many different types of sensors, as well as an overview of the algorithms and tools for data processing and analysis that we have developed and are currently integrating in our system.

Our problem consideration and deployment experiences have shown encouraging signs that fine-grained monitoring for providing services will be possible in the near future. To achieve that one needs high precision sensors for localizing people, preferably without requiring them to

References

1. Bamis A, Savvides A (2009) St: a spatiotemporal filtering language with applications in assisted living. In: PETRA '09: Proceedings of the 2nd international conference on Pervasive technologies related to assistive environments
2. Consolvo S, Roessler P, Shelton B (2004) The CareNet display: lessons learned from an in-home evaluation of an ambient display. UbiComp 2004: Ubiquitous Computing: 6th International Conference, Nottingham, UK, September 7–10, 2004: Proceedings
3. Cook D, Das S (2007) How smart are our environments? An updated look at the state of the art. *Pervasive Mob Comput* 3(2):53–73
4. EATON's Home Heartbeat <http://www.homeheartbeat.com>
5. Enalab (2008) The behaviorscope web (bscopeweb) portal. <http://bscope.eng.yale.edu>
6. European Commission's 6th Framework Programme: service-oriented programmable smart environments for older Europeans (SOPRANO)

7. Gil NM, Hine NA, Arnott JL, Hanson J, Curry RG, Amaral T, Osipovic D (2007) Data visualisation and data mining technology for supporting care for older people. In *Assets '07: proceedings of the 9th international ACM SIGACCESS conference on Computers and accessibility*, New York, NY, USA, pp 139–146
8. Living independently group inc (2008) The quietcare system. <http://www.quietcare.com>
9. Lymberopoulos D, Bamis A, Savvides A (2008) Extracting spatiotemporal human activity patterns in assisted living using a home sensor network. In *PETRA '08: proceedings of the 1st international conference on PErvasive technologies related to assistive environments*, New York, 2008. ACM, pp 1–8
10. Lymberopoulos D, Bamis A, Savvides A (2009) A methodology for extracting temporal properties from sensor network data streams. In *Mobisys '09: proceedings of the 7th international conference on Mobile systems, applications, and services*, New York, NY, USA, 2009. ACM, pp 193–206
11. Lymberopoulos D, Ogale AS, Savvides A, Aloimonos Y (2006) A sensory grammar for inferring behaviors in sensor networks. In *IPSN '06: Proceedings of the 6th international conference on Information processing in sensor networks*, New York, NY, USA, 2006. ACM, pp 251–259
12. Lymberopoulos D, Teixeira T, Savvides A (2007) Detecting patterns for assisted living using sensor networks: a case study. *Sensor Technologies and Applications, 2007. SensorComm 2007. International conference on*, 14–20 October 2007, pp 590–596
13. Lymberopoulos D, Teixeira T, Savvides A (2008) Macroscopic human behavior interpretation using distributed imager and other sensors. *Proc IEEE* 96(10):1657–1677
14. Malan D, Fulford-Jones T, Welsh M, Moulton S (2004) Code-Blue: an ad hoc sensor network infrastructure for emergency medical care. In: *International workshop on wearable and implantable body sensor networks*
15. Pollack M (2005) Intelligent technology for the aging population. *AI Mag* 26(2):9–24
16. Tapia E, Intille S, Larson K (2004) Activity recognition in the home using simple and ubiquitous sensors. *Pervasive computing: second international conference, Pervasive 2004, Linz/Vienna, Austria, April 18–23, 2004: Proceedings*
17. Teixeira T, Culurciello E, Park JH, Lymberopoulos D, Barton-Sweeney A, Savvides A (2006) Address-event imagers for sensor networks: evaluation and modeling. In *IPSN '06: proceedings of the 6th international conference on Information processing in sensor networks*, New York, NY, USA, 2006. ACM, pp 458–466
18. Teixeira T, Savvides A (2007) Lightweight people counting and localizing in indoor spaces using camera sensor nodes. *Distributed smart cameras, 2007. ICDSC '07. First ACM/IEEE international conference on*, 25–28 September 2007, pp 36–43
19. The BehaviorScope project <http://www.eng.yale.edu/enalab/behaviorscope.htm>
20. Wood A, Virone G, Doan T, Cao Q, Selavo L, Wu Y, Fang L, He Z, Lin S, Stankovic J (2006) ALARM-NET: wireless sensor networks for assisted-living and residential monitoring. University of Virginia, Computer Science Department Technical report
21. Yu AS, Bamis A, Lymberopoulos D, Teixeira T, Savvides A (2008) Personalized awareness and safety with mobile phones as sources and sinks. In: *International workshop on urban, community, and social applications of networked sensing systems (UrbanSense08)*