

An Agent-Based Architecture for Sensor Data Collection and Reasoning in Smart Home Environments for Independent Living

Thomas Reichherzer^(✉), Steven Satterfield, Joseph Belitsos, Janusz Chudzynski, and Lamar Watson

Department of Computer Science, University of West Florida, 11000 University Pkwy,
Pensacola, FL 32514, USA

{treichherzer, jchudzynski}@uwf.edu,
{sms23, jrb68, lw28}@students.uwf.edu

Abstract. There has been a tremendous growth in new sensor technology and wireless, handheld computing devices that give rise to new opportunities for smart home applications. With advances in the technology, reduced costs of operation, and the ubiquity of WiFi and cellular networks, the reach and value of smart home technology is growing giving rise to applications of smart home systems to support independent living of the elderly. Such systems must collect and analyze data in the home to recognize unusual activities and alert care givers and/or family members in emergency situations. To address this challenge we developed a flexible and scalable multi-agent system for sensor data collection, integration, processing, and alert management. A prototype system for activity recognition has been built and preliminary tests demonstrate that activities can be successfully recognized based on data captured in the home. As a next step, the sensor network will be deployed in an inpatient residence facility to collect real-world data for evaluating the system's performance and to develop new applications.

Keywords: Agent-based architecture · Sensor network · Activity recognition · Case-based reasoning

1 Introduction

Smart devices are increasingly present in our daily lives and they are projected to be progressively in demand over the next 20 years for applications of healthcare driven in part by the need to improve healthcare services and products for an aging population. Projections show that as the baby-boomer generation reaches age 65 at the rate of 10,000 per day, nearly one in every five Americans will be over the age of 65 by 2029 [1]. Increasing healthcare costs and a growing shortage of professional caregivers as well as a strong desire of elderly people to age in their home are strong predictors of the coming market for smart home technology for elder care [2, 3].

If smart home technology is to gain a wider acceptance for applications in elder care, it must be capable of monitoring the resident around the clock, assisting residents with their daily tasks and alerting family members and caregivers in emergency situations.

To provide such capabilities, the technology must acquire contextual information from sensors in the home and reason without intervention or maintenance to make decisions and provide support. The technology must be extensible allowing newly available sensors to be easily integrated into the smart home while not interfering with other, existing components of the system. Furthermore, the system's reasoning capabilities must be adaptable using machine-learning techniques to be able to change to subtle variations in the behavior patterns of the resident and new emerging needs in the home.

This paper describes the current status and recent achievements in building a Smart Independent Living for Elders (SMILE) home focusing on the sensor technology, system architecture, and revised machine-learning methods, specifically case-based reasoning, based on prior work [4] to improve activity recognition and adaptation capabilities. The objective for the SMILE project is to use off-the-shelf sensors and small computing devices for building wireless sensor networks and applying novel computing algorithms for data processing and analysis to provide support for SMILE residents in the form of reminders or suggestions to complete activities.

The paper is organized as follows: the second and next section describes the sensor network and the agent-based architecture for collecting and analyzing sensor data, providing reasoning and learning support, and generating responses. The third section describes the agents that have been implemented for the prototype system along with a short description of the reasoning algorithm used for activity recognition. The fourth section describes recent experiments and their results to evaluate the performance of the applied methods. The fifth and final section discusses future experiments and directions for the SMILE project.

2 The SMILE Home Multi-agent Architecture

An important design criteria for a smart home system's architecture is its ability to incorporate new technology into the home to meet the needs of the residents without causing disruption to services already provided by the home. Any bug fixes or upgrades to software that improve system reliability and performance must not be disruptive to the system's data collection, processing, or response services. Thus, the SMILE home was designed to be comprised of independent, distributed components modeled as software agents that use multiple databases to record and exchange messages. Each SMILE device, whether a sensor, a widget for visualization in a browser, or an actuator, is managed by a software agent that acts as a wrapper to the technology providing an interface for data collection and control of the device to generate a response of the home.

There are three different types of software agents operating within the SMILE home, each integrated within a three-layer system architecture comprised of a sensor, a middle-ware, and an application layer. Sensor agents operate at the sensor layer to collect raw sensor data from programmable sensors in the home and store them in a sensor database. They record environmental changes such as the light and temperature in a room, the movement of the resident or things, or the status of appliances. Sensor agents may run on small computing devices permanently installed in the home or on mobile devices carried by a resident.

Middleware layer agents process captured sensor data transforming and combining them to describe events in the home and infer a resident’s activities. They can either subscribe to specific events triggered by the database such as the addition of certain environmental data points or poll information of interest from the databases. Agents at the middleware layer use the system’s databases for information sharing which enables the agents to respond to specific types of information they need to process to produce new information for other agents to act upon. There is no centralized control system needed to coordinate information and work flows. Middleware layer agents are run by an application server that also executes multiple Web services to provide secured, controlled access to information stored in the system’s databases to application layer agents including interface agents and all actuators.

The application layer agents provide control services for devices in the home such as turning off and on the light, offer suggestions or warnings to the residents to support activities or prevent accidents, and provide information about the resident’s well being to care takers and family members. The agents may execute on handheld devices or on small computing devices that drive flat screen displays installed in the home to provide opportunities for residents to interact with the SMILE home by accessing recommendations that agents generate or giving feedback to the system.

3 Data Collection, Processing, and Reasoning

As part of the SMILE project a wireless sensor network was built from multiple, low-cost sensors wired to Arduino boards, Raspberry Pi computing devices, and the Nordic nRF24L01+ wireless Radio Frequency (RF24) transceiver modules [5] that connect to the Arduino boards and Raspberry Pis. The RF24 transceivers form together a mesh network for sending data wirelessly from the Arduino boards to the Pis as the sensor base station. Each Arduino board executes a Sensor Data Collector (SDC) software for collecting data from attached sensors; the Raspberry Pi executes the Sensor Data Distribution (SDD) software that receives data packets from one or several SDCs through the RF24 mesh network. The SDD passes received data to sensor agent, also executing on the PI, via internal UDP communication. Figure 1 illustrates the devices and the process for collecting sensor data.




Fig. 1. Collecting sensor data in the home.

The home’s sensor agents execute on a Raspberry Pi or an Android mobile device. Each agent is programmed to collect data from a specific sensor that is given a unique sensor ID. Multiple agents can execute on the same Pi to collect data from different sensors that are within the Pi’s reach of the wireless network. Currently, the home

provides sensor agents for collecting environmental data such as ambient temperature and light, pressure on sofa chair cushions or a bed to detect residents sitting or laying on them, stove top temperatures and infrared sensors to detect when a burner is on and in use, water flow of a sink faucet, and water level in a kitchen sink. These sensors have been chosen to collect sufficient data to infer normal activities such as getting up in the morning and using the stove and the kitchen sink to prepare a meal.

The information and aggregation agents poll the sensor database at close intervals to transform raw sensor data into meaningful descriptions of environmental events as Resource Description Framework (RDF) triples. The agents use an ontology that models entities and entity types in the home such as rooms, appliances, environmental conditions as well as other things in the home and store produced triples in a triple store database. The activity recognition agent uses case-based reasoning methods [6] to infer activities using data from the triple store and a knowledge base that stores descriptions of activities in the home as a set of individual cases. The case representation for each case includes (1) a textual description and a classification of the case, (2) the origin of the case, which can be either adapted or designed, (3) a history of the case, (4) a description of the problem space including temporal and spatial constraints that must be met for a case to match, and (5) a solution space including the recognized activity and risks the activity may pose on the resident or the home. A case's problem space is modeled using a set of weighted features describing events that are relevant to an activity. The weight associated with each feature, a proportional value among all features in the problem description, indicates the strength of relevance. It is assigned by the knowledge engineer but future work will explore techniques for learning them from training data. A history field records how a case has been previously used for recognition of successfully and unsuccessfully completed activities or whether it had been used for generating new cases using adaptation techniques.

To recognize activities, the agent polls the triple store at regular intervals for new events and attempts to map those events onto the features of cases in the case base. This comparison of the events representing the current environmental context with the existing cases in the case base is performed in a two-stage process: an initial surface-level comparison of the events occurring in the home with the RDF triples representing the known patterns in activities so that a set of possible activities can be isolated, followed by a deeper examination of the spatio-temporal aspects of events in the home. In the initial comparison stage, the agent (1) selects new cases from the case base whose features match any of the newly observed events to compile a list of candidate cases, (2) checks if candidate cases match additional features in the cases' problem space descriptions to reach a minimum threshold for deeper examination of the cases, and (3) checks if features no longer match in candidate cases due to changes in events. In a second stage, the agent examines marked cases more closely to find if the temporal and spatial constraints involving case feature match observed events.

During this stage the agent forms expectations on how to complete an activity successfully by examining features not yet matching in candidate cases but expected to match next. If an event occurs that deviates from what is expected by a case, the agent records an expectation failure. It then decides whether cases may be adapted to model new activities not yet recognized to match observed events or whether the observed

events truly represent a failed execution of an activity and warrants intervention by the home. The agent uses annotations associated with constraints to distinguish between adaptable cases and cases representing a serious violation of an activity. Finally, cases representing confirmed activities may be demoted to inactive cases when changes in observed events no longer match features or constraints in the cases.

The interface agents are designed to provide information about the health and well being of the residents and to generate suggestive actions to the residents of the home. Currently, the system provides an information visualization agent operating inside a Web-based dashboard that visualizes the current situation in the home such as temperature and light of individual rooms, status of appliances, and resident location. The agent executes queries on regular intervals to poll environmental information from the sensor and system databases and update graphical widgets integrated into the dashboard. The dashboard makes such information available to trusted individuals. It will be used in upcoming experiments to study patterns of activities in the home.

4 Experiments and Results

A prototype SMILE home has been built and deployed in a test lab environment for initial experimentations. In total, eight sensors (light, temperature, infrared, pressure, water-level sensors), two Raspberry Pis for collecting and distributing sensor data, and ten different iBeacons for indoor localization have been installed. In addition database, application, and Web servers have been deployed on physical servers of the Computer Science department's data center to run middle-layer services. A basic dashboard with widgets for visualization has been implemented and deployed.

To evaluate the introduced algorithm for activity recognition, an experiment was conducted involving several activities in the home. However, as individual components of the system are not yet fully integrated, a synthesized data set of observed events was created and used for the experiment involving five different activities of daily living. The data set was created by a student capturing real-world events that might occur in the home as a result of a series of envisioned activities planned over the course of a day, recording the events as time-stamped RDF triples in the triple store with annotations of the original activity that generated them for evaluation purposes. To make the data set more realistic, noise was added in the form of randomly inserted events describing movements of the resident and objects in the home that are not directly related to the envisioned activities. Finally, a case-base was manually created modeling a total of nine different activities in the home (cooking, managing finances, reading for study, reading for leisure, cleaning the library, cleaning the kitchen, washing dishes, watching television, going to bed) including the five activities that were envisioned by the student.

The algorithm for the activity agent was then executed using the triple store and the knowledge base as described above. Intermediate results of selected and matching cases were recorded. The results were that all nine cases as modeled by the case base were at some point selected as candidate cases of matching activities, six of them were closely examined, and five were recognized as activities matching exactly the activities that were performed during the course of the day. However, the experiment uncovered a

weakness in one of the cases, “cooking on left front burner”. The case did not include a constraint to adequately recognize an expectation failure when the resident leaves the burner on and walks away from the kitchen. Future work will examine methods to build and evaluate cases for completeness before they are added to the case base.

5 Conclusions and Future Work

This paper describes ongoing work of the SMILE project that aims to build an intelligent system inferring activities and providing suggestions to the residents in the home. In collaboration with the Center on Aging at The University of West Florida and Covenant Hospice, the prototype system will be deployed in the near future to an inpatient residence facility. It will be used in a number of already approved human-subject studies in which residents of the facility perform daily activities to record real-world data. The data will be used to (1) evaluate architectural bottlenecks and the accuracy of inferred activities, (2) develop a process to build a case base with critical cases not yet sufficiently covered, and (3) study issues of trust and security in the system. The ultimate goal will be to develop new applications for SMILE residents, care-takers, and family members to support independent living of at-risk older adults.

References

1. Cohn, D., Taylor, P.: Baby Boomers Approach 65 – Glumly. Pew Research Center, December 2010. <http://pewresearch.org/pubs/1834/baby-boomers-old-age-downbeat-pessimism>
2. Harkleroad, J.: Boomers’ desire to ‘age in place’ driving new home trends. Silicon Valley/San Jose Bus. J., 26 February 2006. <http://www.bizjournals.com/sanjose/stories/2006/02/27/focus3.html>
3. El-Basioni, B.M.M., El-Kader, S.M.A., Eissa, H.S.: Designing a local path repair algorithm for directed diffusion protocol. *Egypt. Inf. J.* **13**(3), 155–169 (2012)
4. Satterfield, S., Reichherzer, T., Coffey, J., El-Sheikh, E.: Application of structural case-based reasoning to activity recognition in smart home environments. In: 11th International Conference on Machine Learning and Applications, Vol. 1, pp. 1–6. IEEE (2012). doi:[10.1109/ICMLA.2012.10](https://doi.org/10.1109/ICMLA.2012.10)
5. Nordic Semiconductor, nRF24L01+ Single Chip 2.4 GHz Transceiver Product Specification. Nordic Semiconductor (2007). <http://www.nordicsemi.com/>
6. Aamodt, A., Plaza, E.: Case-based reasoning: foundational issues, methodological variations, and system approaches. *AI Commun.* **7**(1), 39–59 (1994)



<http://www.springer.com/978-3-319-34110-1>

Advances in Artificial Intelligence
29th Canadian Conference on Artificial Intelligence,
Canadian AI 2016, Victoria, BC, Canada, May 31 - June
3, 2016. Proceedings
Khoury, R.; Drummond, C. (Eds.)
2016, XXIV, 338 p. 85 illus., Softcover
ISBN: 978-3-319-34110-1