

A Context-aware System for Ambient Assisted Living

Alessandra De Paola, Pierluca Ferraro, Salvatore Gaglio, Giuseppe Lo Re,
Marco Morana, Marco Ortolani, and Daniele Peri

DIID Department, University of Palermo, Italy.
{firstname.lastname}@unipa.it

Abstract. In the near future, the world's population will be characterized by an increasing average age, and consequently, the number of people requiring for a special household assistance will dramatically rise. In this scenario, smart homes will significantly help users to increase their quality of life, while maintaining a great level of autonomy. This paper presents a system for Ambient Assisted Living (AAL) capable of understanding context and user's behavior by exploiting data gathered by a pervasive sensor network. The knowledge inferred by adopting a Bayesian knowledge extraction approach is exploited to disambiguate the collected observations, making the AAL system able to detect and predict anomalies in user's behavior or health condition, in order to send appropriate alerts to family members and caregivers. Experimental results performed on a simulated smart home prove the effectiveness of the proposed system.

Keywords: Ambient Assisted Living; Multi-sensor data fusion; Dynamic Bayesian Networks; Context awareness; Rule-based Reasoning.

1 Motivations and Related Work

Nowadays, the world's elderly population is rapidly increasing. Recent studies carried out by the World Health Organization report that, by 2020, people aged 60 years and older will outnumber children younger than 5 years, and that there will be more than two billions of elderly people by 2050 [23]. The rising average age of the world's population implies an increase of chronic diseases and requests for specialized assistance.

In such scenario, smart homes might significantly improve elderly and disabled people's quality of life and safety, by helping them live independently and comfortably with the aid of intelligent devices pervasively deployed in the environment. Such vision is in line with the paradigm of Ambient Assisted Living (AAL) [18], a specialization of Ambient Intelligence (AmI) which aims to design intelligent pervasive systems to provide optimal environmental conditions for elders or disabled people.

Understanding user's behavior is one of the key functionalities of these systems, since the detection or prediction of anomalies and dangerous situations can trigger specific actions, e.g., controlling the smart home automation system, sending appropriate alarms to the users or to their family and caregivers [4]. Moreover, the smart home can warn users if they are skipping their usual activities, such as having lunch at the right time or taking medicines, reminding them of the most probable activity to perform in a particular context. This is especially useful for patients suffering from dementia [15], [20]. Obviously, in a smart home, special users will also take advantage

of all usual benefits of AmI systems, which unobtrusively modify the environment to improve users' comfort and satisfy their needs.

To achieve such high-level goals, unobtrusive sensors (e.g., motion and door sensors, or bluetooth beacons to locate users) and actuators can be pervasively deployed in the smart home to monitor and modify ambient properties such as temperature, humidity and lighting. In the specific case of AAL, wearable sensors such as accelerometers and gyroscopes embedded in smart watches, smartphones, or wristbands, can also be used to recognize relevant physical activities, e.g., sitting down or standing up, as well as to automatically detect falls [14]. Moreover, if the user suffers from specific health conditions, continuous monitoring of vital signs can be achieved by exploiting specialized devices [16], [21]. A survey on AAL tools can be found in [18].

The large amount of raw data collected from all these heterogeneous sensors requires an efficient and accurate information fusion process to infer user's activity, current situation and potential anomalies. To this end, multi-sensor data fusion algorithms have been widely adopted [11]; in particular, probabilistic techniques such as Dynamic Bayesian Networks (DBNs) are able to explicitly model the unavoidable uncertainty derived from sensor noise and malfunctioning [6]. DBNs have been described in many works for different applications, such as user presence detection [7], target tracking [24], and fire detection [1].

As demonstrated by a large body of research on the topic [2, 9, 17], raw measures acquired by ad hoc sensor can be enriched by considering context information that can be exploited to further increase the inference accuracy. This is especially useful in AAL scenarios in order to distinguish among ambiguous situations. For example, an increase in body temperature and blood pressure might be normal if the user is exercising, or might indicate an anomaly if the user is laying down on the bed in distress [19]. Given the frequency of ambiguous situations in a real world setting, and the importance of minimizing false positives and false negatives, accurately determining the user's context is paramount to develop a functional AAL system which can effectively assist elderly people [19].

In this work we present a multi-layered architecture for a complete AAL system, capable of understanding the current situation and detecting anomalies in users' behavior and health condition, as well as planning the best course of actions to help them avoiding dangerous situations and living safely in their own home. Our system is based on a context-aware, multi-sensor data fusion module able to infer users' activities and describe the environment through high level concepts. Context information is correlated with vital signs data coming from specialized sensors to unambiguously determine if the user is in distress. A rule-based reasoner has been designed to plan the best sequence of actions to perform, so as to ensure the safety and well-being of the monitored users.

The remainder of this paper is organized as follows. Section 2 introduces the multi-layered architecture of the proposed system, focusing on the activity recognition and reasoning modules. Section 3 describes how mixing sensor data and context attributes can be exploited to increase the accuracy of the system. Section 4 presents the experimental setting and the evaluation results. Finally, Section 5 draws our conclusions with directions for future work.

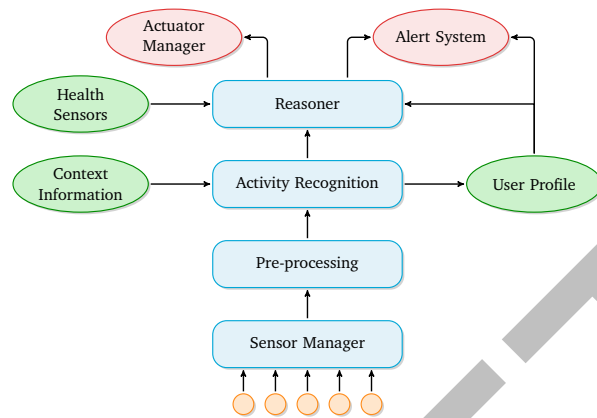


Fig. 1: Multi-layered architecture of the proposed system.

2 Multi-layer Architecture

The AAL system we propose here is based on a multi-layer architecture (Fig. 1) designed to collect raw data from heterogeneous sensors and aggregate them with increasing levels of abstraction, so as to unambiguously understand the current situation of the user (e.g., normal, abnormal, emergency). The inferred high level knowledge is exploited by a symbolic reasoner that plans the best sequence of actions to perform.

Modules at the lowest levels of the architecture are responsible for managing the sensory infrastructure. The *Sensor Manager* module sends control messages to sensors and collects raw data needed by the higher levels of the architecture. Data stream is analyzed by the *Pre-processing* module which detects outliers, filtering out unreliable data, in order to improve the accuracy of the system.

One of the most relevant component of the AAL system is the *Activity Recognition* module, that adopts a probabilistic approach based on Dynamic Bayesian Networks (DBN) to deal with uncertain sensory data. Low-level context information (e.g., time and location) and sensory data are used to recognize the activity performed by the user, so producing high-level context information related to the currently performed activities. The knowledge inferred by the Activity Recognition module is used to build a user profile, which represents the historical user behavior through a DBN.

At the topmost level of the AAL system, the rule-based symbolic *Reasoner* infers the user's condition by merging data coming from specialized health sensors, the user's activity as detected by the Activity Recognition module, and the user's profile historical data. The Reasoner is also responsible for planning the best sequence of actions to perform in order to guarantee the user's safety and well-being, even through the generation of automatic and manual alerts. When the user's condition is considered anomalous, the Reasoner triggers the *Alert system*, which sends appropriate alerts to users and their caregivers. Moreover, the user is able to manually ask for a suggestion about the next activity to perform; such functionality is particularly useful for users affected by dementia. The Reasoner also triggers the *Actuator Manager*, which coordinates sensors and

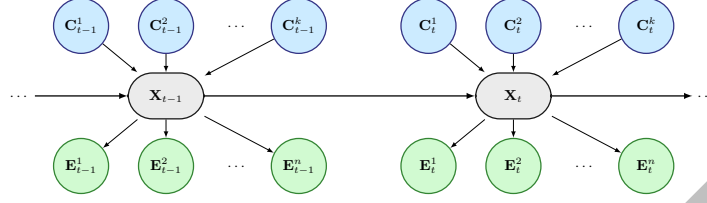


Fig. 2: Structure of the DBN used for inferring the activity performed by the user.

actuators, modifying the state of the environment by controlling heating, ventilation, air conditioning (HVAC) and lighting systems, in line with the paradigm of Aml.

The following sections provide an in-depth description of the most complex modules, i.e., the Activity Recognition module and the Reasoner.

2.1 Activity Recognition Module

The Activity Recognition module is one of the core modules of the system. It performs a multi-sensor data-fusion process by exploiting a Dynamic Bayesian Network (DBN) to handle dynamic phenomena and to deal with the unavoidable inaccuracy of sensors. DBNs are a generalization of Hidden Markov Models (HMMs), which allows the adoption of general topologies. They model the time flow as a sequence of slices, where each slice represents the state of the observed phenomena at a given time. Each slice is composed of an arbitrary number of state nodes, which represent the knowledge to be inferred, and a set of evidence nodes, which represent known information.

The DBN adopted by the Activity Recognition module is shown in Fig. 2. The hidden variable X_t represents the observed phenomena, i.e., the activity performed by the user at any time slice t . To update its belief, the system exploits evidence nodes, which are composed of a set of sensor nodes, $\mathbf{E}_t = (E_t^1, \dots, E_t^n)$, and a set of context nodes, $\mathbf{C}_t = (C_t^1, \dots, C_t^k)$ in the time slice t . Context nodes provide further information by exploiting readily available context attributes, such as time and location. The role of context information in our system will be analyzed in detail in Section 3.

The proposed DBN is fully characterized by its structure and its *sensor* and *state transition* models. The *sensor model*, defined as the probability distribution $P(\mathbf{E}_t|X_t)$, expresses the correlation between sensor readings and the value of the state variable. The *state transition model*, defined as $P(X_t|X_{t-1}, \mathbf{C}_t)$, is the probability distribution indicating how likely is each value of the state variable, given its previous value and all available context information. The DBN performs inference by maintaining and updating its *belief* about a specific system state, x_t , in the time slice t , as follows:

$$Bel(x_t) = P(x_t|\mathbf{E}_{1:t}, \mathbf{C}_{1:t}). \quad (1)$$

Eq. (1) can be expressed by the following recursive equation:

$$Bel(x_t) = \eta \cdot \prod_{e_t^i} P(e_t^i|x_t) \cdot \sum_{x_{t-1}} P(x_t|x_{t-1}, \mathbf{C}_t) \cdot Bel(x_{t-1}), \quad (2)$$

where η is a normalizing constant. The derivation of Eq. (2) is given in [6] and is omitted here due to space limitations. According to Eq. (2), the Activity Recognition module can update its belief about the user's activity by storing only the last two slices of the DBN. Time and space complexity of the inference process are thus independent of the sequence length. The conditional probability tables which define sensor and state transition models can be learned in different ways, depending on the available training data, e.g., sample statistics, Expectation Maximization algorithm or gradient ascent [12].

2.2 Rule-based Reasoner

The classic monitor-analyze-plan-execute cycle implemented by most AMI systems is particularly useful for designing adaptive systems that operate in dynamic scenarios [10]. Such paradigm requires continuous monitoring of the surrounding environment, and allows the system to quickly react in case of emergency.

The proposed *Reasoner* adopts a rule-based approach to infer the user's health and comfort condition, on the basis of information received from lower-levels modules, and consequently plan the sequence of actions to perform in the current situation in order to satisfy user's needs. The Reasoner has been implemented using the rule-based inference engine Jess (Java Expert System Shell) [8]. Jess exploits a pattern matching algorithm to query a knowledge base and act according to logical rules, expressed with a LISP-like programming language. Each Jess rule has the "if <conditions> then <actions>" form, and it is executed when all of the specified conditions are satisfied.

One of the most relevant properties at the basis of the Reasoner behavior is named *user-condition*. It describes the inferred condition of the user, and can take three possible values, namely *normal*, *anomaly*, and *emergency*. Depending on the value of the *user-condition* property, the system changes its alert level, and consequently, dynamically modifies the type and amount of sensory information used in the data fusion process, together with the recipients, the frequency and the level of detail of alert messages. A *normal* condition indicates that the user is healthy and no anomalies or emergencies were detected; the system enters a power saving mode, turning off unused sensors and focusing on discreetly monitoring the user, guaranteeing his comfort, and sending sporadic alert messages. An *anomaly* occurs when the readings from specialized health sensors are slightly out of the norm, or when the user is skipping his usual activities (e.g., taking medicines at the right time). In such cases, the system requires increased accuracy to ascertain the gravity of the situation and, if the anomaly persists, it starts sending alerts. Finally, the system enters in *emergency* mode when it detects a serious or prolonged anomaly, by analyzing both vital signs and information from the *Activity Recognition* module.

In the following, a small subset of the rules exploited by the system are described to illustrate how the system infers the *user-condition* property and modifies its own behavior dynamically. The first rule puts the system into the *anomaly* condition if the user's blood pressure is out of the norm and the user is laying on the bed:

defrule anomalyHighBloodPressure:

if (*blood pressure* is "high") **and** (*activity* is "laying on bed") **then**
user-condition ← "anomaly"

It is worth noticing that the same blood pressure reading is normal if the user is exercising, as shown by the following rule:

defrule normalHighBloodPressure:

if (*blood pressure* is “high”) and (*activity* is “exercising”) then
 user-condition ← “normal”

In this case the high blood pressure can be explained by the stress of doing physical exercise. The system correctly disambiguates the sensory readings and infers that the *user-condition* is *normal*.

Finally, the following rule is triggered when the system detects a medical *emergency* and the user is laying on the floor in distress:

defrule emergencyHeartRate:

if (*heart rate* is “very high”) and (*activity* is “laying on the floor”) then
 user-condition ← “emergency”
 send alerts to family members, caregivers and medical staff

The system recognizes the seriousness of the emergency and immediately sends appropriate alerts to the involved people and medical staff.

It is worth noticing that the user can also force the system to enter in *anomaly* state, by explicitly asking for assistance, or in *emergency* state by indicating a state of distress.

3 Context-awareness

In the architecture proposed here, context information is exploited at multiple levels. The multi-level nature of the system allows and encourages reuse of information acquired or inferred by lower-level modules. In such architecture, higher-level modules exploit context information in a transparent way, totally decoupled from the implementation of lower levels. For example, knowledge about the activity performed by the user can be exploited by high level components, even neglecting the algorithm adopted by the *Activity Recognition* module, that can be considered as a smart *virtual sensor* responsible for perceiving the activity performed by the user. Such abstraction mechanism can be applied to all modules below the *Reasoner*, as depicted in Fig. 3, which suggests an alternative view of the proposed architecture, highlighting the data flow between the modules of the AAL system. To be more specific, knowledge flows from low to high levels, starting from raw data acquired by physical sensors, up to a high-level description of the user’s condition and the environment.

In the application scenario of AAL, physical sensors representing the input of the knowledge extraction process can be classified in three categories: ambient sensors, personal (or inertial) sensors, and specialized health sensors. Ambient sensors, such as lighting, temperature, motion and door sensors, are pervasively deployed in the smart home and constitute the main input of the *Activity Recognition* module. Information coming from these sensors can be fused to infer the user’s location with room-level granularity, so feeding a *virtual location* sensor which produces part of the context information exploited by higher-level modules.

Personal sensors, e.g., accelerometers and gyroscopes, might be embedded in smartphones, smart watches or smart wristbands, and thus follow users wherever they go. The

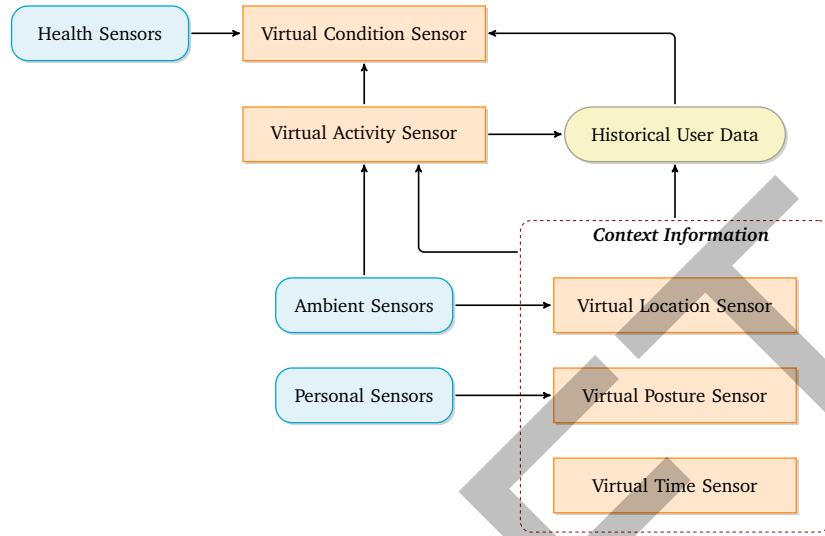


Fig. 3: Data flow through modules of the proposed architecture.

fusion of their readings can be seen as a *virtual posture* sensor capable of detecting if the user is sitting, standing up, or laying on the floor or on the bed.

Besides location and posture, the set of context information is enriched with simple time-related information (e.g., the time of day, the month, or the day).

Context information and ambient sensory readings are the input of the *virtual activity* sensor, whose output is in turn combined with historical user data to infer the *user context*, which is a high-level description of what the user is doing and why. An example of user context is “*the user is correctly taking her morning medicines after breakfast*”. Inferring such detailed information is possible by correctly recognizing the sequence of actions performed by the user (e.g., having breakfast followed by taking medicines) and comparing it to the user’s daily schedule, which may include taking specific medicines after breakfast. Then, to effectively determine the *user condition*, the output of the *virtual activity* sensor and the inferred *user condition* are fused with data coming from specialized health sensors, which measure vital signs such as heart rate and blood pressure. Finally, this symbolic description is leveraged by the Reasoner module, as described in Section 2.2.

4 Experimental Analysis

Experimental Setting. The proposed system is evaluated in a simulated smart home, where several sensor devices are pervasively deployed in all rooms. Sensor traces were obtained from the Aruba dataset of the Washington State University’s CASAS Smart Home Project [3].

Such data was collected over a period of seven months in a smart home, with eight rooms and a single resident, by means of 31 motion sensors, 5 temperature sensors, and

Table 1: Diversity index, average duration (hh:mm:ss), and number of time slices associated with each activity, based on the CASAS dataset [3].

Activity	Diversity Index	Avg. duration	# Time slices
Bed to toilet	0.713	00:02:30	912
Eating	1.340	00:10:30	5,366
Enter home	1.296	00:00:30	173
Housekeeping	1.942	00:20:30	1,373
Leave home	1.156	00:00:30	173
Meal preparation	0.997	00:08:00	25,246
Outside	0.734	01:50:00	89,498
Relax	1.064	00:36:00	197,072
Resperate	0.749	00:09:00	110
Sleeping	0.615	04:03:30	192,433
Wash dishes	0.853	00:07:30	978
Work	0.800	00:18:00	6,010
Other	1.895	00:12:30	114,244

3 door sensors. The original data has been preprocessed by grouping sensor events into time windows of 30 seconds, counting the number of activations of each sensor in each window. Temperature data has been discarded, given its low correlation with the user’s activities, as previously shown in [6].

The Aruba dataset is annotated with 11 activities of daily living, namely *eating*, *meal preparation*, *wash dishes*, *sleeping*, *bed to toilet*, *housekeeping*, *relax*, *work*, *enter home*, *leave home*, and *resperate* (i.e., using a device for the treatment of high blood pressure). Two other activities have been added to the dataset as a preprocessing step, i.e., *outside* and *other*. The former covers the periods when the user is not at home (intervals between *leave home* and *enter home*), while the latter groups sensor events not belonging to known activities, as proposed by [13].

Impact of Location-Related Context Information. The first set of experiments presented here aims to prove the relevance of location-related context information, which is extrapolated from ambient sensors and grouped in a *Virtual Location Sensor*, as described in Section 3. In particular, we extensively studied different features of the most common daily living activities performed in a smart home in the application scenario of AAL. Our findings are summarized in Table 1, which reports the diversity index, average duration, and number of time slices associated with each activity in the considered setting. The diversity index is a measure that reflects how evenly the activities are carried out in different functional areas of the smart home. In particular, we divided the smart house in rooms, assigning each sensor to one of them. Then, we analyzed the ground truth of each activity to identify those performed exclusively in a single room, and those carried out in different locations.

We adopted the Shannon-Wiener index, i.e., the classical Shannon entropy [22], to evaluate the diversity index of activities. Accordingly to such metric, an activity that is performed in a well-defined location, such as *sleeping*, has a low diversity index, while

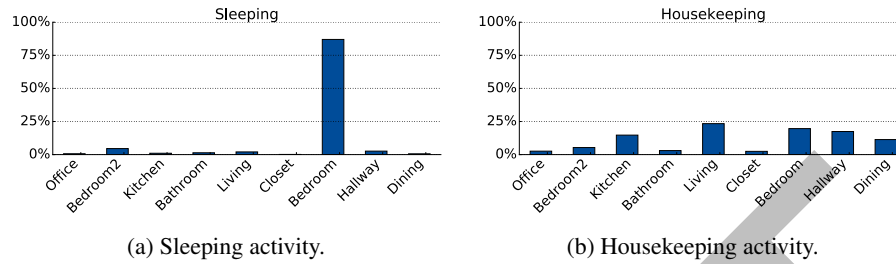


Fig. 4: Rooms where the activities with the lowest (a) and highest (b) *diversity index* are performed.

activities carried out in numerous rooms of the smart house, such as *housekeeping* and *other*, have a high diversity index.

To better illustrate this aspect, we selected the activity with the lowest diversity index (i.e., *sleeping*), and the one with the highest index (i.e., *housekeeping*). Fig. 4 shows the rooms where the *sleeping* and *housekeeping* activities were performed. As expected, *housekeeping* is carried on in almost all rooms, whilst *sleeping* is concentrated in the bedroom. As it turns out, many activities present low diversity indices (in fact, more than half of them are less than 1) being performed in very specific locations, like *sleeping*. Thus, adding context information related to the current user’s position greatly improve the accuracy of the system, as will be shown in the following.

Location information is easily obtainable in a reliable way by exploiting cheap motion sensors. Simply knowing the room where each sensor is deployed is enough to infer the room the user is in as the one where more motion sensors are active in the current time window. If the smart home has a single resident, as is often the case in AAL scenarios and in the chosen dataset as well, this simple counting operation is enough to guarantee a highly reliable virtual location sensor with room-level granularity, as confirmed by the experiments we performed. If the smart home is occupied by multiple residents, more advanced techniques involving the use of bluetooth beacons or RFID chips can be exploited to identify and locate each user.

Impact of Other Context Information. Evaluating the accuracy of the *Activity Recognition* module is paramount to assess the whole system behavior, since if the system fails to correctly infer the activity performed by the user, higher level modules are unable to work reliably. We verified that context information is useful to improve the accuracy of such module.

As expected, it can be observed that the activities with a high diversity index are difficult to be correctly recognized, while regular activities, such as *sleeping*, are classified more accurately. Similar remarks apply to the average duration and the number of occurrences reported in Table 1: it is more difficult to classify activities that are short, or that occur very rarely.

Table 1 also highlights that most activities have similar average durations (with some obvious exceptions), and thus this attribute is not useful to better disambiguate

between them. However, this kind of information can be extremely valuable for energy optimization systems, as described in [6].

As regards the *Virtual Time Sensor*, we found out that the most impactful attribute is the period of day (i.e., morning, afternoon, evening, or night), which greatly influences the activities performed by users. Surprisingly, other time-related attributes, such as month or day of week, are less impactful, and sometimes even harmful for accuracy.

Activity Recognition Evaluation. Finally, the whole system has been evaluated. It is worth noticing that, since the Reasoner has a deterministic behavior, its accuracy is directly dependent on the accuracy of the Activity Recognition module.

In order to evaluate the inference accuracy, we adopted the following metric:

$$Acc = \frac{TP + TN}{TP + TN + FP + FN}, \quad (3)$$

where TP , TN , FP , and FN are, respectively, the true positives, true negatives, false positives and false negatives.

We also calculated the *precision* (or positive predictive value) and *recall* (or sensitivity) to provide more details on fidelity and completeness of the system, as follows:

$$precision = \frac{TP}{TP + FP}, \quad recall = \frac{TP}{TP + FN}. \quad (4)$$

Finally, the harmonic mean of precision and recall, named *F-score*, is computed as well with the following formula:

$$F-score = 2 \cdot \frac{precision \cdot recall}{precision + recall}. \quad (5)$$

All experiments have been performed by using the cross-validation method. The dataset was partitioned into ten parts, nine of which were used for learning the DBN parameters, whilst the tenth was used for evaluating the system; this process was repeated ten times by changing the test set and averaging the results.

Fig. 5 shows the accuracy, average uncertainty, precision, recall, and F-score obtained by exploiting both the virtual time and location sensors. Choosing the right combination of context attributes is fundamental to improve the accuracy of activity recognition systems, as it can modify such results by up to 12% [5]. For this reason, great attention has been paid to selecting optimal context attributes in the application scenario of AAL. Fig. 5 shows that the system achieves a remarkable average accuracy of 0.805, by choosing the best combination of context information available. This result is especially good because it includes the *other* activity which is particularly difficult to recognize. If we don't consider this ambiguous activity, the accuracy of the system reaches 0.897, with an increase of almost ten percent. Finally, the uncertainty of the system, defined as the classic Shannon entropy, is satisfactorily low, with a reported value of 0.294 over a maximum of 3.7.

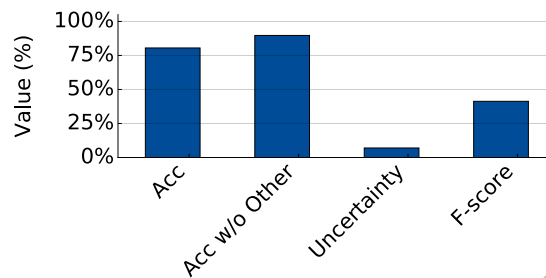


Fig. 5: Accuracy, accuracy ignoring the *other* activity, normalized uncertainty, and F-score of the proposed system.

5 Conclusions

In this paper, we proposed a context-aware system for Ambient Assisted Living, which aims to improve elderly and disabled people’s quality of life by helping them living safely and independently in their own home. The proposed multi-layered architecture allows the system to understand the current situation and detect anomalies and emergencies in user’s health conditions and behavior, reminding them if they skip an important activity in their daily routine, and possibly alerting family members and caregivers.

A context-aware multi-sensor data fusion module infers users’ activities with remarkable accuracy, as shown in the experimental section. Context information is exploited at various levels to improve the inference accuracy and disambiguate between different situations. At the highest level, a rule-based inference engine plans the best course of actions when the user needs help, by assessing the gravity of the situation and sending alerts to specialized people that can rapidly assist the user. Even though all the parts of the proposed architecture have been individually tested, an overall evaluation of the whole AAL system, as well as a comparison with other state of the art approaches, is still in progress and will be presented in a future work. Currently, the rules exploited by the *Reasoner* are manually defined by a domain expert. In the future, we are interested in automatically learn such rules, making the system even more adaptive.

Finally, as suggested in the experimental section, multi-user scenarios would be both extremely challenging and interesting to study as the next step in developing a complete AAL system. The biggest challenge in this area comes from the fact that users’ activities can influence each other, and thus personalized users’ models are needed to effectively handle multiple users living in the same smart apartment.

References

1. Cheng, N., Wu, Q.: A decision-making method for fire detection data fusion based on Bayesian approach. In: Proc. 4th Int’l Conf. on Digital Manufacturing and Automation (ICDMA). pp. 21–23. IEEE (2013)
2. Cho, K., Hwang, I., Kang, S., Kim, B., Lee, J., Lee, S., Park, S., Song, J., Rhee, Y.: HiCon: a hierarchical context monitoring and composition framework for next-generation context-aware services. *IEEE Network* 22(4), 34–42 (2008)

3. Cook, D.J.: Learning setting-generalized activity models for smart spaces. *IEEE Intelligent Systems* 2010(99), 1 (2010)
4. Cook, D.J., Youngblood, M., Das, S.K.: A multi-agent approach to controlling a smart environment. In: *Designing smart homes*, pp. 165–182. Springer (2006)
5. De Paola, A., Ferraro, P., Gaglio, S., Lo Re, G.: Context-awareness for multi-sensor data fusion in smart environments. In: *AI* IA 2016 Advances in Artificial Intelligence*, pp. 377–391. Springer (2016)
6. De Paola, A., Ferraro, P., Gaglio, S., Lo Re, G., Das, S.: An adaptive bayesian system for context-aware data fusion in smart environments. *IEEE Trans. on Mobile Computing* (2016)
7. De Paola, A., La Cascia, M., Lo Re, G., Morana, M., Ortolani, M.: User detection through multi-sensor fusion in an AmI scenario. In: *Proc. 15th Int'l Conf. on Information Fusion (FUSION)*. pp. 2502–2509. IEEE (2012)
8. Friedman, E.: *Jess in action: rule-based systems in Java*. Manning Publications Co. (2003)
9. Huebscher, M.C., McCann, J.A.: Adaptive middleware for context-aware applications in smart-homes. In: *Proc. 2nd Workshop on Middleware for Pervasive and Ad-Hoc Computing*. pp. 111–116. ACM (2004)
10. Kephart, J., Chess, D.: The vision of autonomic computing. *Computer* 36(1), 41–50 (2003)
11. Khaleghi, B., Khamis, A., Karray, F.O., Razavi, S.N.: Multisensor data fusion: a review of the state-of-the-art. *Information Fusion* 14(1), 28–44 (2013)
12. Koller, D., Friedman, N.: *Probabilistic graphical models: principles and techniques*. MIT press (2009)
13. Krishnan, N.C., Cook, D.J.: Activity recognition on streaming sensor data. *Pervasive and Mobile Computing* (2012)
14. Lombardi, A., Ferri, M., Rescio, G., Grassi, M., Malcovati, P.: Wearable wireless accelerometer with embedded fall-detection logic for multi-sensor ambient assisted living applications. In: *Sensors, 2009 IEEE*. pp. 1967–1970. IEEE (2009)
15. Lotfi, A., Langensiepen, C., Mahmoud, S.M., Akhlaghinia, M.J.: Smart homes for the elderly dementia sufferers: identification and prediction of abnormal behaviour. *Journal of ambient intelligence and humanized computing* 3(3), 205–218 (2012)
16. Ni, Q., García Hernando, A.B., de la Cruz, I.P.: The elderly's independent living in smart homes: A characterization of activities and sensing infrastructure survey to facilitate services development. *Sensors* 15(5), 11312–11362 (2015)
17. Padovitz, A., Loke, S.W., Zaslavsky, A., Burg, B., Bartolini, C.: An approach to data fusion for context awareness. In: *Modeling and Using Context*, pp. 353–367. Springer (2005)
18. Rashidi, P., Mihailidis, A.: A survey on ambient-assisted living tools for older adults. *IEEE Journal of Biomedical and Health Informatics* 17(3), 579–590 (2013)
19. Roy, N., Das, S.K., Julien, C.: Resolving and mediating ambiguous contexts in pervasive environments. *Smart Healthcare Applications and Services: Developments and Practices* pp. 122–147 (2011)
20. Roy, N., Pallapa, G., Das, S.K.: A middleware framework for ambiguous context mediation in smart healthcare application. In: *Proc. 3rd IEEE Int'l Conf. on Wireless and Mobile Computing, Networking and Communications (WiMOB)*. pp. 72–79. IEEE (2007)
21. Sardini, E., Serpelloni, M.: T-shirt for vital parameter monitoring. In: *Sensors*, pp. 201–205. Springer (2014)
22. Shannon, C.E.: A mathematical theory of communication. *ACM SIGMOBILE Mobile Computing and Communications Review* 5(1), 3–55 (2001)
23. Suzman, R., Beard, J.R., Boerma, T., Chatterji, S.: Health in an ageing world - what do we know? *The Lancet* 385(9967), 484–486 (2015)
24. Zhang, Y., Ji, Q.: Active and dynamic information fusion for multisensor systems with dynamic Bayesian networks. *IEEE Trans. Systems, Man, and Cybernetics, Part B: Cybernetics* 36(2), 467–472 (2006)