# *Context-Awareness for Multi-Sensor Data Fusion in Smart Environments*

Article

Accepted version

A. De Paola, P. Ferraro, S. Gaglio, G. Lo Re

AI*IA 2016 Advances in Artificial Intelligence. AI*IA 2016. Lecture Notes in Computer Science, vol 10037

DOI: 10.1007/978-3-319-49130-1_28



Publisher: Springer

# Context-Awareness for Multi-Sensor Data Fusion in Smart Environments

Alessandra De Paola*, Pierluca Ferraro, Salvatore Gaglio, and Giuseppe Lo Re

DICGIM Department, University of Palermo, Italy.
{alessandra.depaola, pierluca.ferraro, salvatore.gaglio, giuseppe.lore}@unipa.it

**Abstract.** Multi-sensor data fusion is extensively used to merge data collected by heterogeneous sensors deployed in smart environments. However, data coming from sensors are often noisy and inaccurate, and thus probabilistic techniques, such as Dynamic Bayesian Networks, are often adopted to explicitly model the noise and uncertainty of data.
This work proposes to improve the accuracy of probabilistic inference systems by including context information, and proves the suitability of such an approach in the application scenario of user activity recognition in a smart home environment. However, the selection of the most convenient set of context information to be considered is not a trivial task. To this end, we carried out an extensive experimental evaluation which shows that choosing the right combination of context information is fundamental to maximize the inference accuracy.

**Keywords:** Multi-sensor data fusion; Dynamic Bayesian Networks; Context awareness

## 1 Motivations and Related Work

Nowadays, users expect end-applications to provide useful context-aware services, by exploiting the increasing number of sensors deployed in smart environments and smartphones [1]. To this end, pervasive computing applications need to accurately infer the current context, by efficiently processing large amounts of raw sensory data.

For this purpose, multi-sensor data fusion is extensively used to combine data collected by heterogeneous sensors [2]. However, since sensor data are often noisy and inaccurate, probabilistic techniques are widely adopted to explicitly model the noise and uncertainty of raw data, as described in [3]. In particular, Dynamic Bayesian Networks (DBNs) [4] take into consideration the past belief of the system, in addition to data coming from sensors, and allow to handle the dynamicity of the observed phenomena. Many works leverage DBNs to perform adaptive data fusion for different applications, such as fire detection [5], target tracking [6], and user presence detection [7, 8]. A detailed survey on multi-sensor data fusion can be found in [9].

Many systems presented in the literature exploit context information to improve the inference accuracy and reduce the uncertainty of unreliable sensor data [10]. Multi-attribute utility theory is exploited in [11] for modeling and merging context attributes,

* Corresponding Author

with the goal of achieving situation awareness. The authors of [12] propose a context aggregation framework that can recognize context information of various scale (i.e., personal, local, city-wide, and global) and combine it hierarchically. Moreover, various frameworks use context information to reduce unnecessary communications among wireless sensors, thus reducing their energy consumption [13, 14].

We propose a context-aware multi-sensor data fusion system to infer high-level context information about the world, that includes low-level context information in order to refine the inference process. The output of the inference process can be further exploited by higher level reasoning modules to derive new knowledge, in a multi-layered architecture that aims to provide a symbolic description of the environment.

We demonstrate the effectiveness of the proposed approach in an Ambient Intelligence (AmI) [15] scenario, whose goal is to create smart environments which satisfy users' needs, by exploiting pervasive sensors and actuators that surround users, with a low level of intrusiveness [16]. To meet such requirement, many AmI designers prefer to use low-cost and low-impact devices, possibly already deployed in the environment, rather than developing ad-hoc sensors to specifically monitor the features of interest, and thus the collected data are usually only partially related to observed phenomena [8].

In the field of AmI, a key challenge is recognizing users' activities [17, 18]. Various approaches have been proposed in the literature, depending on the kind of activities to classify. For example, to recognize activities of daily living (e.g., sleeping, working, eating), wireless sensors are often unobtrusively deployed in smart environments, so as not to bother users [3]. Conversely, inertial sensors, such as those commonly found in smartphones, are better suited to recognize activities that involve physical movements, e.g., sitting down, walking, and running [19, 20].

We focus on user activity recognition in a smart home environment, and exploit context information at different levels. The inference of context information, as a high-level description of the users' activity, is the main goal of the system. Moreover, basic context attributes, such as time-related and location-related information, are used to refine the inference process. Such basic context attributes can be reliably and easily sensed, and thus do not increase the uncertainty of the system.

Unlike other works presented in the literature, we advocate that it is not always convenient to blindly include all available context information in the data fusion process. On the contrary, as we demonstrate in the experimental section, choosing the right combination of context information is fundamental to maximize the inference accuracy. To this end, we propose to exploit only context attributes which are readily available and easy to measure in a reliable way, so as not to increase the uncertainty of the system. Moreover, we prove that choosing the right combination of context information is fundamental to maximize the inference accuracy, especially when only few sensors are available. In such cases, our results show that exploiting context information improves the accuracy of the system by almost $13\%$.

The remainder of this paper is organized as follows. Section 2 describes the multi-layered architecture of the proposed system, focusing on the context-aware DBN that performs the inference. Section 3 discusses the context information that can be exploited to increase the accuracy of the system. Section 4 presents the experimental

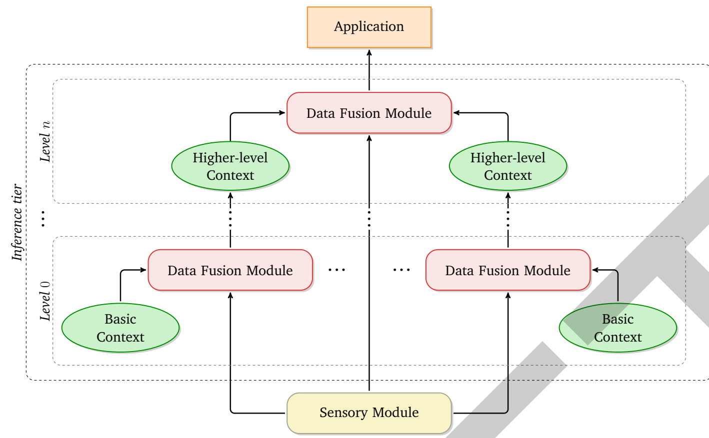


Fig. 1: Multi-layered architecture of the context-aware data fusion system.

setting and the results of our analysis. Finally, Section 5 draws our conclusions and proposes directions for future work.

## 2 Multi-layer Architecture

This paper proposes a novel approach to multi-sensor data fusion for intelligent systems based on the use of pervasive sensors. One of the main features of the system is its capability of dealing with inaccurate and noisy data coming from sensory devices. In particular, the use of probabilistic techniques allows our system to merge information coming from multiple sensors by explicitly modeling the noise and uncertainty of data [9].

Fig. 1 shows the multi-layered architecture of the system. At the lowest tier, the *Sensory* module perceives the world through the pervasive sensory infrastructure. The inference tier is composed of multiple levels: at each level, one or more *Data Fusion* modules exploit context attributes coming from lower levels to perform probabilistic inference on the pre-processed sensory data, fusing them to infer new context information which provides a higher level description of the environment. The process of knowledge abstraction continues until the context information requested by the top-level application is inferred.

In this work, we will focus on a single *Data Fusion* module, and on the impact that context information has on its inference accuracy. A more accurate description of the *Data Fusion* module is presented in the following section.

### 2.1 Data Fusion Module

The proposed data fusion system is based on a DBN, which models the observed phenomena taking into account the past state of the world besides current sensory readings. DBNs are a specialization of Bayesian Networks that guarantee a great flexibility in

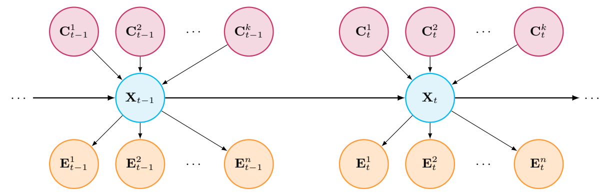


Fig. 2: Structure of the Dynamic Bayesian Network (DBN) used for the inference.

model expressiveness [21]. They pose no restrictions on conditional probability distributions, differently from Kalman filters [22], and allow for more general topologies than Hidden Markov Models [23]. A DBN is partitioned in temporal slices, where each slice represents the state of the world in a given moment, besides the evidences representing the observable manifestation of the hidden state of the world. Each slice of a DBN can have any number of state variables and evidence variables.

Fig. 2 shows the structure of the DBN we designed. Our goal is to infer the state of the world, in the form of a given feature of interest, on the basis of a set of sensory readings, represented by the evidence nodes $\boldsymbol{E}_t = (E_t^1, \ldots, E_t^n)$ at any time slice $t$. Differently from prior work, we also exploit a set of context information, represented by the evidence nodes $\boldsymbol{C}_t = (C_t^1, \ldots, C_t^k)$ in the time slice $t$. We will analyze in detail the choice of which context information to use in Section 3.

To fully characterize the DBN, it is necessary to define the *sensor model* and the *state transition model* [4]. The probability distribution $P(\boldsymbol{E}_t|X_t)$ expresses how sensory readings are affected by the state variable, and is named *sensor model*. The *state transition model*, defined as $P(X_t|X_{t-1}, \boldsymbol{C}_t)$, represents the probability that the state variable takes a certain value, given its previous value and the current context information.

The *belief* of the system about a specific value of the state variable at time $t$ is defined as:

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

By following a procedure analogous to that adopted in [24] for deriving the equation of Bayes filters, it is possible to express Eq. (1) in the following recursive formulation:

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

where $\eta$ is a normalizing constant. Using Eq. 2, we only need to store the last two slices of the DBN, and thus the time and space required for updating the belief do not increase over time. Calculating the belief for a single $x_t$ has a computational complexity of $O(n + m)$, where $n$ is the number of sensor nodes and $m$ is the number of possible values of the state variable. The overall complexity of computing $Bel(x_t)$ for all values of $X_t$ is therefore $O(m^2 + m \cdot n)$.

In order to populate the conditional probability tables of the DBN, several different methods can be adopted, depending on the training set. In a fully labeled dataset, we can compute sample statistics for each node. Otherwise, if the values of one or more of the variables are missing for some of the training records, we can adopt the Expectation Maximization (EM) algorithm or some form of gradient ascent [25].

## 3 Context-awareness

The role of context in our system is twofold. First, our main goal is the inference of context information, intended as a high-level description of the surrounding world. In particular, as described in Section 1, we are interested in recognizing the activities performed by users in a smart home environment, which in turn will enable higher-level applications to provide to users the most appropriate services.

Low-level context information, such as time and location, can be exploited by our data fusion system to improve the accuracy of reasoning by refining the inference process, as demonstrated by many context-aware data fusion systems proposed by researchers over the years [11, 26].

However, using too many context attributes can actually be detrimental to the inference accuracy, as will be demonstrated in Section 4.3, and increases the computational burden of the system, especially in the training phase. Thus, it is important to analyze the possible context information and select only the most informative attributes, which may vary depending on the application scenario.

We identify some principles that should drive the selection of context attributes. First of all, context information should be readily available in all situations, regardless of the sensors used. Therefore, we suggest to discard information provided by users manually, together with context attributes which are difficult to sense or that cannot be sensed directly and reliably, thus introducing new elements of uncertainty in the system.

The authors of [27] provide a widely accepted definition of context, which identifies the *primary* categories of contextual information, i.e., identity, activity, location, and time. Identity and activity are high level attributes, while location and time are low level attributes. Thus, according to the principles stated above, we will focus on location-related and time-related context information, analyzing the possible benefits they can provide to the system, and validating our intuitions in the experimental section.

Time-related context information is used by most context-aware systems in literature, since it is very easy to obtain (i.e., it is sufficient to check the current date and time). For activity recognition systems, in particular, time-related context information provides remarkable improvements to the accuracy [28]. First of all, intuitively, activities performed by users may vary a lot in different periods of day: for example, sleeping is the most probable activity during the night, and many users have lunch and dinner at regular time each day. Thus, exploiting this context attribute should improve the accuracy of the system, with almost no drawbacks. However, the number of periods in which a day is divided can influence the performance of the system, as we will demonstrate in Section 4.3. Both too coarse-grained periods (e.g., intervals of 12 hours) and too fine-grained ones (e.g., intervals of 1 minute) do not convey much information; hence, finding the best granularity is very important.

Similarly, activities performed by users might be influenced by the current day of the week and, to a lesser extent, by the month of the year. However, we expect the activities of users to be less correlated to these context attributes, with respect to the period of day. For example, it is possible that users will behave differently during weekends, but it is unlikely that activities will change much among the other days. We defer further considerations regarding the day of the week and month of the year to the experimental section. Other time-related context information, such as the timezone, might be interesting for different scenarios, but are irrelevant for our case study of activity recognition in a smart house.

As regards location-related context information, we focus primarily on the position of users, leaving to future work an analysis on how to exploit the position of objects to improve the awareness of the system about users' surroundings. In the case of a smart home, with no strong assumption on the kind of sensors used, we propose to exploit user location information with a room-level granularity. Regardless of the sensors used, estimating the position of users with this level of detail is required to correctly inferring their activities.

However, a system that relies primarily on location-related context information will encounter difficulties in recognizing certain activities. Intuitively, this can be explained by considering that some activities are performed in well-defined locations (e.g., sleeping in the bedroom), and therefore are well recognized using this kind on information, while other activities are more irregular (e.g., housekeeping, which may be carried out in all rooms of the smart home), and more heterogeneous context information should be exploited to recognize them with higher accuracy.

## 4 Experimental Analysis

In order to evaluate the possible contribute of different context information to the data fusion process, we test the performance of the proposed system while varying the type and granularity of context information.

### 4.1 Simulation Setting

We evaluated our system in a simulated smart home, pervaded by several sensor devices, as proposed by [29]. Sensory traces and corresponding user activities were obtained from the Aruba dataset of the CASAS Smart Home Project [3], at Washington State University. This dataset contains annotated data collected in a smart apartment with a single resident, over a period of seven months. Events are generated by 31 motion sensors, 3 door sensors, and 5 temperature sensors, deployed in 8 rooms (5 sensors per room on average).

We partitioned the sequence of sensor events into time windows of 30 seconds, counting how many times each sensor was activated during each slice. We noticed a low correlation between temperature readings and the activity performed by the user, and thus we decided to discard this information.

The Aruba dataset considers eleven activities of daily living (ADLs), i.e., *Bed to Toilet*, *Eating*, *Enter Home*, *Housekeeping*, *Leave Home*, *Meal Preparation*, *Relax*, *Resperate*[1], *Sleeping*, *Wash Dishes*, and *Work*. We added a new activity, named *Outside*,

[1] Resperate is a device used for the treatment of high blood pressure.

that takes into consideration the periods of time when the user is not at home, i.e., the intervals between *Leave Home* and *Enter Home*.

We also added another activity, named *Other*, which groups all the sensor events that do not match any of the known activities. We think it is essential to detect this activity class accurately in a real world scenario, since nearly 20% of the sensor events in the dataset considered here belong to the *Other* class. However, considering the heterogeneity of the activities grouped by this class, it is very challenging to recognize it with good accuracy, and many approaches in the literature ignore it altogether, relying on a static list of known activities, as noted in [17].

We used the cross validation method to evaluate the system, dividing the dataset into ten parts. For each test, nine parts were used for learning the CPTs (Conditional Probability Tables) of the DBN, and the tenth was used for the test. This process was then repeated changing the test set ten times and averaging the results.

After the pre-processing phase, the dataset consisted of $633\,468$ sensor events. Each test of the cross validation used $570\,121$ sensor events as training cases, and $63\,347$ sensor events as test cases. All experiments have been performed on a workstation equipped with an Intel® Core$^{\text{TM}}$ i5-3470 CPU (4 cores, 3.20 GHz, 4 GB RAM). The training phase required $4\,914$ ms on average.

### 4.2 Performance Metrics

We adopted the average accuracy as metric to evaluate the performance of the activity recognition systems, defined as:

$$Acc = \frac{TP + TN}{TP + TN + FP + FN}, \tag{3}$$

where *TP*, *TN*, *FP*, and *FN* are, respectively, the true positives, true negatives, false positives and false negatives. However, accuracy alone is not sufficient to evaluate different approaches, since data are skewed towards the most probable activities. In fact, activities such as *Sleeping* and *Relax* account for a large number of time slices, while others like *Resperate* and *Leave Home* or *Enter Home* are much rarer and shorter. For this reason, we adopted additional metrics to provide a more detailed analysis of the performance of the systems.

To measure the uncertainty of the probabilistic reasoning performed by the systems, we used an index based on the classic definition of Shannon entropy [30]. We also calculated the average cross-entropy error function, which is defined as follows:

$$CE = -\frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{M} y_{ij} \log p_{ij}, \tag{4}$$

where $N$ is the number of timesteps, $M$ is the number of activity classes, and $y_{ij}$ and $p_{ij}$ are, respectively, the ground truth and the predicted probability for the $j^{th}$ activity class at time $i$. The cross-entropy error is an information-theoretic measure of accuracy that incorporates the idea of probabilistic confidence, measuring the cross-entropy between the distribution of true labels and the prediction of the system. This kind of error becomes extremely large (i.e., $+\infty$ in the extreme case) if the system is over-confident about a wrong prediction, and it is thus useful to evaluate the accuracy of the belief with a fine granularity. Finally, we determined the *precision* (positive predictive value), as fidelity measure, and the *recall* (sensitivity), for measuring completeness, which are defined as follows:

$$precision = \frac{TP}{TP + FP}, \qquad recall \quad = \frac{TP}{TP + FN}. \tag{5}$$

Precision and recall, in turn, are used to calculate the *F-score*, defined as the harmonic mean of precision and recall, as follows:

$$\textit{F-score} = 2 \cdot \frac{precision \cdot recall}{precision + recall}. \tag{6}$$

### 4.3 Experimental Results

**Time-related context information**

The first set of experiments we present is a detailed analysis on the importance of some time-related context attribute, i.e., *period of day*, *day of week*, and *month*.

We will begin by studying the performance of the system when changing the granularity of the *period of day* node. Fig. 3a shows the accuracy, uncertainty, F-score and cross-entropy error of a system exploiting the *period of day* node, as a function of the number of periods in which a day is divided, starting from a single period (i.e., a single interval of 24 hours) up to a maximum of 48 periods (i.e., 48 intervals of 30 minutes). We notice an increment of the accuracy and F-score when increasing the granularity up to 6 periods (i.e., intervals of 4 hours). Likewise, uncertainty and cross-entropy error are very low using this granularity. However, if we divide the day in more than 6 periods, we observe a steady decrease of the F-score, as well as an increase of uncertainty, whilst accuracy and cross-entropy remain unchanged. Thus, we can conclude that increasing the time granularity is beneficial only up to a point; going further only adds to the noise, resulting in a system that performs worse with no added benefits.

Our experimental results show that it is possible to improve accuracy and F-score of the system even more by dividing the day manually in four periods, namely *morning* (8AM - 12PM), *afternoon* (12PM - 8PM), *evening* (8PM - 11PM) and *night* (11PM - 8AM). This way, the periods closely follow the phases of the day when the type of activities performed by typical users changes, as shown in Fig. 3b. As the figure points out, the user's behavior changes remarkably during the day. For example, the *Housekeeping* and *Wash Dishes* activities are much more probable during morning or afternoon, and it seems the user works prevalently on afternoons. As expected, activities such as *Sleeping* and *Bed to Toilet* take place mainly at night. However, some activities, such as *Other*, show less variance during the day, and are thus more difficult to identify. Results show that this granularity yields the best accuracy, F-score and uncertainty ($0.793$, $0.416$, and $0.231$, respectively), and one of the lowest cross-entropy errors, i.e., $1.858$.

In order to evaluate the effect of context information concerning the day of week and the month, we analyzed the frequency of the user's activities, during the week (Fig. 3c) and among different months (Fig. 3d). It is worth noting that the user's behavior is pretty regular during the week, including weekends. The only exception appears

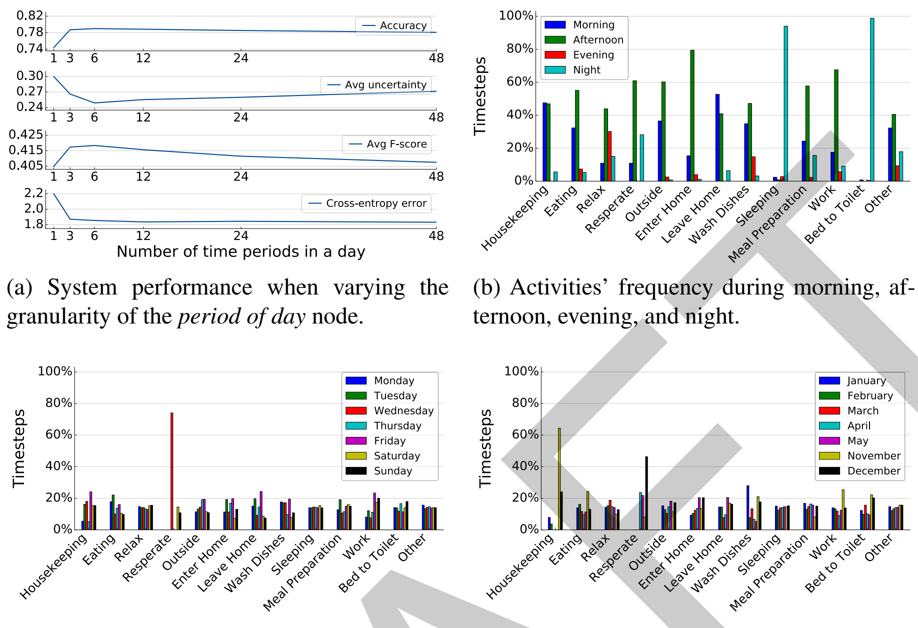


(a) System performance when varying the granularity of the *period of day* node.

(b) Activities' frequency during morning, afternoon, evening, and night.

(c) Activities' frequency in different days.

(d) Activities' frequency in different months.

Fig. 3: Analysis on the importance of time-related context information.

to be the *Resperate* activity; however, this is a fairly rare activity, thus its weight when determining the accuracy of the system is limited. Even among different months, the activities are quite regular (only *Housekeeping* and *Resperate* show a remarkable variance). By analyzing such results, it is easy to predict the different impact of the *period of day* node with respect to the *day of week* and *month* ones.

In order to verify our analysis, we compared eight systems that exploit different combinations of such context information, as reported in Table 1. The difference in accuracy between the best and worst combination of context nodes is more than $10\%$. We can observe that four systems out of the five with highest accuracy exploit the *period of day* node. Moreover, these systems show high F-scores and low uncertainty and cross-entropy errors. As expected, the accuracy of all systems improves significantly if the *Other* activity class is ignored, increasing by about 10% on the average. Surprisingly, the system which includes all three context nodes performs worse than the one which excludes them. This can be explained by the interference of the *month* and *day of week* nodes. In fact, the system that exploits only these two context nodes is the worst according to all the metrics. Conversely, the system that performs better is the one which uses only the *period of day* node. Activities are too regular during the week and among months, and therefore the usefulness of the *day of week* and *month* nodes is limited. Thus, at a first glance, it appears that the *day of week* and *month* nodes are not needed to improve the performance of the data fusion system, and can, in fact, be detrimental.

Table 1: Average accuracy (Acc), uncertainty, cross-entropy error (CE), and F-score of the analyzed systems, sorted by accuracy in descending order.

| Period of Day | Day of Week | Month | Acc | Acc w/o *Other* | Uncertainty | CE | F-score |
|---|---|---|---|---|---|---|---|
| ✓ | – | – | 0.793 | 0.889 | 0.231 | 1.858 | 0.416 |
| ✓ | ✓ | – | 0.779 | 0.874 | 0.245 | 1.950 | 0.400 |
| ✓ | – | ✓ | 0.778 | 0.876 | 0.280 | 1.815 | 0.385 |
| – | – | – | 0.760 | 0.853 | 0.282 | 2.146 | 0.403 |
| ✓ | ✓ | ✓ | 0.739 | 0.833 | 0.373 | 1.911 | 0.366 |
| – | ✓ | – | 0.734 | 0.826 | 0.347 | 2.232 | 0.390 |
| – | – | ✓ | 0.714 | 0.800 | 0.429 | 2.222 | 0.363 |
| – | ✓ | ✓ | 0.690 | 0.772 | 0.562 | 2.347 | 0.349 |

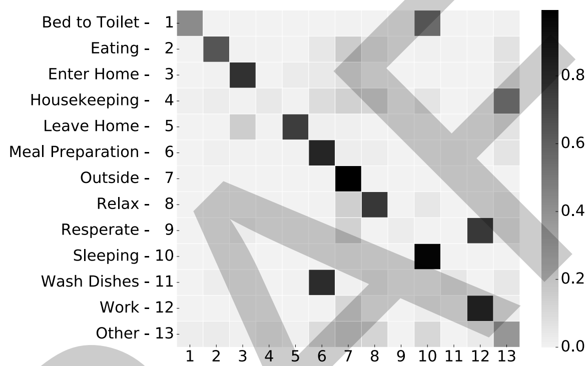


Fig. 4: Confusion matrix of the baseline data fusion system.

The system which uses only the *period of day* node will be considered as baseline for comparison with other systems in next experiments. To provide a more detailed analysis of its performances, its confusion matrix, row-wise normalized, is presented in Fig. 4. Each cell $C_{ij}$ represents the number of instances of class $i$ predicted to be in class $j$ by the system. Therefore, diagonal entries correspond to true positives, and non diagonal entries correspond to classification errors. To explain why some activities are more difficult to recognize than others, in the following we will analyze the location in which each activity is carried out.

**Location-related context information**

Intuitively, we can hypothesize that some activities are performed in well-defined locations, and therefore are well recognized using only motion sensors, while other activities are more irregular. Furthermore, we supposed that some activities are performed mainly in the same rooms (and roughly in the same time periods), such as *Wash Dishes* and *Meal Preparation*. To verify these hypotheses, we divided the smart house in rooms, and measured the variability of the association between activities and rooms, through the *diversity index*, defined as the classical Shannon entropy [30]. Fig. 5a summarizes the *diversity index* of the activities, which indicates how they are carried out in different rooms. Activities performed in a well-defined location have a low diversity index,

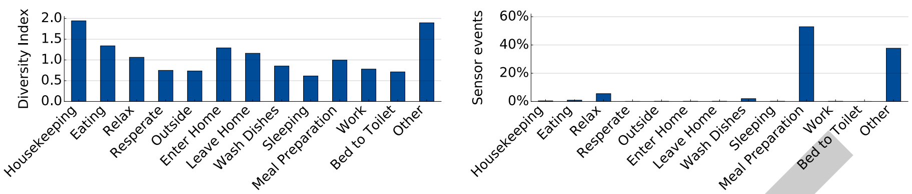


(a) *Diversity index* of each activity. (b) Frequency of activities in the kitchen.

Fig. 5: (a) *Diversity index* of the activities, and (b) frequency of activities performed in the kitchen.

while activities carried out different rooms exhibit a high diversity index. As expected, activities which are difficult to recognize correctly, such as *Housekeeping* and *Other*, exhibit the highest diversity indices. On the other hand, activities that are easier to classify accurately, such as *Sleeping*, have low *diversity indices*. The *Wash Dishes* activity seems to contradict this statement, since it sports a low *diversity index*, but is often misclassified. However, this activity only takes place in the kitchen, since almost $80\%$ of sensor events associated with it comes from sensors deployed there. This accounts for its low *diversity index*, but, as Fig. 5b shows, there are much more probable activities taking place in the same room, such as *Meal Preparation* ($52.9\%$ of sensor events) and *Other* ($37.7\%$ of sensor events); even *Relax* is a more probable activity than *Wash Dishes*, in the kitchen. Therefore, it is understandable that a system which relies mostly on motion sensors will have a hard time identifying this kind of activity. To overcome this problem, we can exploit the information associated to the duration of each activity.

**Duration of activities**

We observed that some activities exhibit a much longer average duration than others. For example, *Sleeping* has an average duration of about 4 hours, while *Eating* generally takes about 10 minutes. Thus, it is intuitive that making use of this kind of context information should be beneficial to the system.

In order to verify the usefulness of such information, we tested a system with an additional context node exploiting the duration of activities, and compared it to our baseline system. Surprisingly, the resulting accuracy was lower than the baseline system by about $2\%$, and the other metrics were unchanged or slightly worse. A closer look at the data reveals that only a couple of activities (i.e., *Sleeping* and *Outside*) have average durations longer than one hour. Most of the other activities have durations similar to each other, generally between 10 and 30 minutes. It seems that this type of context information fails to help the system if we can exploit enough data coming from the sensory devices.

However, when performing data fusion, it might not always be efficient to sample all available sensors. On the contrary, it may be useful to activate only a subset of sensors, depending on the application scenario. For instance, if the sensory infrastructure is composed of devices with limited energy resources, the use of a subset of devices might increase the lifetime of the whole network.

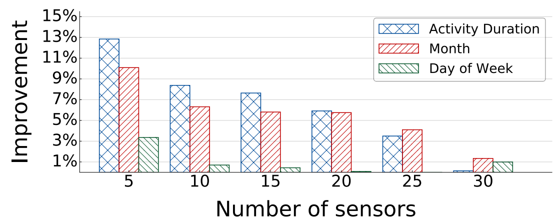


Fig. 6: Improvement of inference accuracy when exploiting context information with different number of sensors.

For this reason, we repeated the comparison experiments using only a subset of sensors, discarding the rest of the data. As expected, in these conditions context information proved to be much more valuable. Using only 10 sensors (out of 34), the accuracy of the baseline system is $65.87\%$, while exploiting duration information results in an accuracy of $74.25\%$, with a significant improvement of $8.38\%$. As it turns out, the same is true for other context information as well.

Fig. 6 shows the improvement in accuracy of systems exploiting *activity duration*, *month* and *day of week*, with respect to the baseline system (i.e., the one exploiting only the *period of day*) as a function of the number of sensors used. It can be noted that, in the extreme case of using only 5 sensors, exploiting the *activity duration* improves the accuracy of the system by almost $13\%$. Conversely, the benefits of using context information decrease when there is enough data coming from the sensory devices. The same holds true for the *month* node, whilst the improvement when using the *day of week* is negligible even with few sensors.

## 5 Conclusions

In this paper, we have proposed a multi-sensor data fusion system that aims to improve the accuracy of probabilistic inference by including context information in the fusion process. The key idea is that context information can be involved at different levels of the reasoning process. Basic context attributes can contribute to improve inference accuracy, as demonstrated in the experimental evaluation. At the same time, the context information inferred as result of the data fusion constitutes a high-level description of the environment, and can be exploited by other reasoning engines to better support top-level applications that provide context-aware services to users.

We have demonstrated the suitability of such approach in the application scenario of user activity recognition in a smart home environment. The experimental results have confirmed that choosing the right combination of context information is fundamental to maximize the inference accuracy, especially when only few sensors are available, and that exploiting the best context information set greatly improves the accuracy of activity recognition systems.

As future work, we are interested in studying how to further use context information to dynamically reconfigure the sensory infrastructure, by sampling a subset of sensors in order to minimize energy consumption, whilst maintaining a high degree of inference accuracy.

Moreover, in this paper, we focused on a scenario involving a single user in a smart apartment, and we will study multi-user scenarios in the future. However, recognizing activities performed by multiple users is really challenging, since users can influence each other. Several studies demonstrated that using personalized models for each user dramatically improve systems performance [31], but learning personalized models is computational expensive, and it is thus necessary to find a good trade-off between accuracy and computational efficience. Finally, we are interested in considering training and test data coming from different smart environments, so as to verify the generalization potential of the system.

## References


1. Lo Re, G., Morana, M., Ortolani, M.: Improving user experience via motion sensors in an ambient intelligence scenario. In: Pervasive and Embedded Computing and Communication Systems (PECCS), 2013. (2013) 29–34
2. De Paola, A., La Cascia, M., Lo Re, G., Morana, M., Ortolani, M.: Mimicking biological mechanisms for sensory information fusion. Biologically Inspired Cognitive Architectures **3** (2013) 27–38
3. Cook, D.J.: Learning setting-generalized activity models for smart spaces. IEEE Intelligent Systems **2010**(99) (2010) 1
4. Murphy, K.P.: Dynamic Bayesian Networks: representation, inference and learning. PhD thesis, University of California, Berkeley (2002)
5. 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), IEEE (2013) 21–23
6. 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) (2006) 467–472
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), IEEE (2012) 2502–2509
8. De Paola, A., Gaglio, S., Lo Re, G., Ortolani, M.: Multi-sensor fusion through adaptive Bayesian networks. In: AI*IA 2011: Artificial Intelligence Around Man and Beyond. Springer (2011) 360–371
9. 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) (2013) 28–44
10. 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, ACM (2004) 111–116
11. Padovitz, A., Loke, S.W., Zaslavsky, A., Burg, B., Bartolini, C.: An approach to data fusion for context awareness. In: Modeling and Using Context. Springer (2005) 353–367
12. 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) (2008) 34–42
13. Kang, S., Lee, J., Jang, H., Lee, Y., Park, S., Song, J.: A scalable and energy-efficient context monitoring framework for mobile personal sensor networks. IEEE Trans. Mobile Computing **9**(5) (2010) 686–702

14. Nath, S.: ACE: exploiting correlation for energy-efficient and continuous context sensing. In: Proc. 10th ACM Int'l Conf. on Mobile systems, applications, and services (MobiSys), ACM (2012) 29–42
15. Cook, D., Augusto, J., Jakkula, V.: Ambient Intelligence: technologies, applications, and opportunities. Pervasive and Mobile Computing **5**(4) (2009) 277–298
16. De Paola, A., Farruggia, A., Gaglio, S., Lo Re, G., Ortolani, M.: Exploiting the human factor in a wsn-based system for ambient intelligence. In: Complex, Intelligent and Software Intensive Systems, 2009. CISIS '09. International Conference on. (2009) 748–753
17. Krishnan, N.C., Cook, D.J.: Activity recognition on streaming sensor data. Pervasive and Mobile Computing (2012)
18. Gaglio, S., Lo Re, G., Morana, M.: Human activity recognition process using 3-d posture data. Human-Machine Systems, IEEE Transactions on **45**(5) (2015) 586–597
19. Gao, L., Bourke, A., Nelson, J.: Evaluation of accelerometer based multi-sensor versus single-sensor activity recognition systems. Medical Engineering & Physics **36**(6) (2014) 779–785
20. Kwapisz, J.R., Weiss, G.M., Moore, S.A.: Activity recognition using cell phone accelerometers. ACM SigKDD Explorations Newsletter **12**(2) (2011) 74–82
21. De Paola, A., Gagliano, L.: Design of an adaptive Bayesian system for sensor data fusion. In: Advances onto the Internet of Things. Springer (2014) 61–76
22. Meinhold, R.J., Singpurwalla, N.D.: Understanding the Kalman filter. The American Statistician **37**(2) (1983) 123–127
23. Sanchez, D., Tentori, M., Favela, J.: Hidden Markov Models for activity recognition in Ambient Intelligence environments. In: Proc. 8th Mexican Int'l Conf. on Current Trends in Computer Science, IEEE (2007) 33–40
24. Thrun, S., Burgard, W., Fox, D.: Probabilistic robotics. MIT press (2005)
25. Koller, D., Friedman, N.: Probabilistic graphical models: principles and techniques. MIT press (2009)
26. Nimier, V.: Introducing contextual information in multisensor tracking algorithms. In: Advances in Intelligent Computing (IPMU). Springer (1995) 595–604
27. Dey, A.K., Abowd, G.D., Salber, D.: A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. Human-Computer Interaction **16**(2) (2001) 97–166
28. Perera, C., Zaslavsky, A., Christen, P., Georgakopoulos, D.: Context aware computing for the Internet of Things: A survey. Communications Surveys & Tutorials, IEEE **16**(1) (2014) 414–454
29. Lalomia, A., Lo Re, G., Ortolani, M.: A hybrid framework for soft real-time WSN simulation. In: Proc. 13th IEEE/ACM Int'l Symp. Distributed Simulation and Real Time Applications (DS-RT). (2009) 201–207
30. Shannon, C.E.: A mathematical theory of communication. ACM SIGMOBILE Mobile Computing and Communications Review **5**(1) (2001) 3–55
31. Weiss, G.M., Lockhart, J.W.: The impact of personalization on smartphone-based activity recognition. In: AAAI Workshop on Activity Context Representation: Techniques and Languages. (2012)