Mobile Netw Appl (2007) 12:155–171
DOI 10.1007/s11036-007-0013-5

Activity Recognition for Context-aware Hospital Applications: Issues and Opportunities for the Deployment of Pervasive Networks

Jesus Favela · Monica Tentori · Luis A. Castro · Victor M. Gonzalez · Elisa B. Moran · Ana I. Martínez-García

Published online: 13 July 2007

Abstract Hospitals are convenient settings for the deployment of context-aware applications. The information needs of hospital workers are highly dependent on contextual variables, such as location, role and activity. While some of these parameters can be easily determined, others, such as activity are much more complex to estimate. This paper describes an approach to estimate the activity being performed by hospital workers. The approach is based on information gathered from a workplace study conducted in a hospital, in which 196 h of detailed observation of hospital workers was recorded. Contextual information, such as the location of hospital workers, artifacts being used, the people with whom they collaborate and the time of the day, is used to train a back propagation neural network to estimate hospital workers activities. The activities estimated include clinical case assessment, patient care, preparation, information management, coordination and classes and certification. The results indicate that the user activity can be correctly estimated 75% of the time (on average) which is good enough for several applications. We discuss how these results can be used in the design of activity-aware applications, arguing that recent advances in pervasive and networking technologies hold great promises for the deployment of such applications.

Keywords activity estimation · context-aware computing · hospital activities · neural networks

J. Favela (✉) · M. Tentori · L. A. Castro · E. B. Moran · A. I. Martínez-García
Computer Science Department, CICESE,
Ensenada, Mexico
e-mail: favela@cicese.mx

M. Tentori
e-mail: mtentori@cicese.mx

L. A. Castro
e-mail: luis.castro@acm.org

E. B. Moran
e-mail: elmoran@cicese.mx

A. I. Martínez-García
e-mail: martinea@cicese.mx

V. M. Gonzalez
School of Informatics, University of Manchester,
Manchester, UK
e-mail: vmgonz@manchester.ac.uk

1 Introduction

Hospitals are good candidates for the introduction of pervasive technology [4, 6]. Clinical environments are filled with increasingly complex technology, including computers and sensors, where patient care requires coordination and collaboration among specialists; and the working staff is highly mobile and technology savvy. Indeed, some elements of pervasive computing are gradually being introduced in hospitals. These range from wireless networks, PDAs [9], RFID tags for patient tracking [27], voice-activated communication devices [32], and sensors for patient monitoring [29].

One of the challenges of hospital work is the management of large amounts of information, including patient records, medical guides, and scientific papers used for evidence-based medicine [21]. The information needs of hospital workers are highly dependent on contextual information such as location, role, time of day, and activity. For instance, for a nurse attending a patient, the document more relevant might be the patient's chart, while for a physician it might be the medical health record.

This has motivated the development of context-aware applications that adapt to changes in the environment to better assist hospital workers [3, 25]. These applications focus mostly on supporting intra-hospital communication and information access based on user location and role. In this regard, considerable work has been done in the development of robust approaches to location estimation for in-door working environments [17]. Although, other contextual variables such as role and time of day can be easily determined, estimating the activity being performed is more complex.

In this work we present an approach for the automatic estimation of the activity being performed by hospital workers. This approach is based on the use of a neural network trained to map from contextual information (e.g., location, artifacts being used) to user activity (e.g., clinical case assessment). The classifier is trained and evaluated with data captured from close to 200 h of detailed observation and documentation of hospital workers.

Activity information could be relevant for a diverse number of hospital applications, such as deciding whom to call for help or facilitating access to relevant patient information. For instance, the Vocera communication system uses a voice-controlled badge to enable mobile users to communicate over the wireless network currently being used in some hospitals [32]. The Vocera system enables a physician to place a call to “a nurse in the emergency unit”. If the user’s activity could be accurately estimated, a system such as this one would be able to decide which of the nurses to call based on their perceived availability. Similarly, considering that hospital workers, and nurses in particular, spend a considerable amount of time documenting their work, if the system were aware of the nurse’s activity, for instance that she is administering medication to the patient at a particular bed, this information could be automatically captured into the system, simplifying this laborious task.

Despite all the benefits of activity-aware applications in hospitals, the development of this type of systems faces important challenges, such as the design and deployment of networks of sensors needed to monitor contextual variables relevant to estimate users’ activities. Recent advances in wireless communications and electronics have enabled the development of low-cost, low-power, multifunctional sensor nodes that are small in size and communicate information in short distances. These tiny sensor nodes, which consist of sensing, data processing, and communicating components, have undoubtedly brought in a new era of pervasive computing with ubiquitous network connectivity. Despite the ample gamut of practical, useful application these pervasive sensors could afford, several issues need to be addressed for the efficient operation of this technology in real working settings.

The rest of the paper is organized as follows: In Section 2 we describe a user study conducted in a hospital to determine the activities performed by hospital staff as well as to gather the data used to train and evaluate the activity classifier. Section 3 describes the architecture and training of the neural network that estimates the activity of users, as well as the contextual variables used as input for the classifier. The results obtained are presented and discussed in Section 4. In Section 5, we present some implications for the design of activity-aware applications for hospital work. Section 6 discusses some of the issues that must be solved for the deployment of networks of pervasive sensors required to estimate hospital workers’ activities. In section VI, we describe the previous work related to activity estimation and how it compares with the approach being proposed. Finally, Section 7 presents our conclusions and directions for future work.

2 Activities performed by hospital staff

We conducted a series of workplace studies in a mid-size public hospital in the city of Ensenada, Mexico. We studied the Internal Medicine area where more than 70% of the patients are attended. A preliminary study was conducted in this area, aimed at revealing the time hospital workers spend performing different activities; the distance they move, the places they move to and the reason of doing it; the people with whom they collaborate more often and the artifacts they use in support of their work [23]. The main contribution of this study consists of the characterization of mobile work and the information management practices that hospital workers engage in. Then we conducted a more detailed study to gather additional data. The time of each action was recorded on paper as it occurred, annotating details regarding the nature of the actions, artifacts used, content of conversations, and physical location of individuals. All time stamps were recorded to the second with as much precision as it was possible. As all records were done by hand at the field site, they were later transcribed and integrated into observation reports to facilitate its analysis and the computation of statistics. Finally, data were grouped into different classes (i.e., activities of users) which were later used on the training of the neural network.

Contextual information, such as the location of hospital workers, artifacts being used, the people with whom they collaborate and the time of the day, can be used to infer the user’s activity if one is familiar enough with the hospital’s rhythms of work. Temporal patterns or “rhythms” are used by hospital workers to coordinate their activities and contribute to the regular temporal organization of work in the hospital [31]. These rhythms are said to be used by hospital workers to infer the activity conducted by colleagues. For instance, one can infer that a physician is in the ward round just by knowing that he and some medical interns are at 11 A.M. in a patient room.

The study was conducted over a period of eight months, where five medical interns, five nurses and five physicians were shadowed for two complete working shifts and interviewed. These roles of informants were selected because they experiment frequent task switching and are responsible for patient care. At the beginning of the study, each informant was introduced to it and asked to participate voluntarily. Figure 1 shows pictures of the informants while working during a day of observation. The total time of detailed observation was about 196 h and 46 min, with an average time of observation per informant of 13 h and 7 min. To enhance our understanding of their practices and to clarify issues, we conducted also 1-h interviews with each informant after the period of observation.

2.1 Scenario: a typical day of a medical intern

To illustrate the different activities and "rhythms" [31] experienced by hospital workers we present a scenario that describes a typical working day of one of the medical interns observed. Figure 2 depicts the activity switching described in the scenario:

The medical interns meet at 7 A.M. with the attending physician at the internal medicine office, where they briefly discuss the night's events described by the intern who worked the night shift. After the discussion, the interns gather the information related to the patients assigned to them and place it in each patient's bedroom. They walk down to the laboratory to gather the laboratory results of the patients, and attach them to the medical record. Later, the medical interns meet at the internal medicine office and for one or 2 h, they listen to a colleague's assessment of a particular medical issue. After that, they go to the bed wards where along with the physician they conduct the ward round. They discuss each patient's clinical case consulting the patient's medical record and their studies. Finally, once the medical intern has finished the round, which occasionally lasts until 1–2 P.M., the rest of the shift is spent mostly doing paperwork at the internal medicine office.

Figure 1 Hospital workers during one of the observation days. **a** A nurse preparing medicines, **b** a nurse documenting clinical information of a patient and **c** medical interns with the attending physician discussing a clinical case sharing the medical record

To understand the medical behaviors experienced by those observed, we conducted a qualitative analysis of the data collected. The particular qualitative approach we followed was inspired by the techniques to derive grounded theory originally proposed by Strauss and Corbin [33]. For our particular case, the qualitative technique involved continuous comparative analysis of the information collected including interview transcripts, personal notes and documents. As a result of this analysis, we developed a coding scheme that describes the activities performed by hospital workers. This coding scheme was cross-analyzed by the team of researchers and some hospital workers to validate and refine it.

2.2 Hospital workers' activities

To clarify each type of activity, we next describe them in more detail including the artifacts used, the people involved and the location where it normally takes place:

Some of the activities identified are carried out with the aim of providing quality of care. For instance, *clinical case assessment* refers to the actions performed by hospital workers to make a clinical decision to assess the evolution of a patient; by examining clinical evidence (e.g., medical records); by discussing a patient diagnoses with colleagues and by consulting reference material (e.g., pharmacological equivalences). While this activity is focused on making a clinical decision, others, such as *patient care*, are focused on providing primary care for patients. This involves the actions conducted with the aim of examining the patient's health, monitoring vital signs, administering medication, advising the patient, and executing specialized procedures. These activities are performed in front of the patient, where interns and physicians jointly execute them, whereas nurses generally perform them alone.

Some of the activities identified could be either executed with the aim of conducting either a main or supporting function. The activities with a support function are coordination, clinical evidence management, and classes and certification. *Coordination* refers to the set of actions performed by hospital workers to manage personnel and to supervise the quality of care provided to patients. *Clinical evidence management* refers to the actions performed by hospital workers to formalize the notes taken on the move creating medical evidence in the form of medical notes or nurse sheets. Finally, *classes and certification* involves assisting to sessions where a group of hospital workers (i.e., specialized personnel, nurses or specialists),

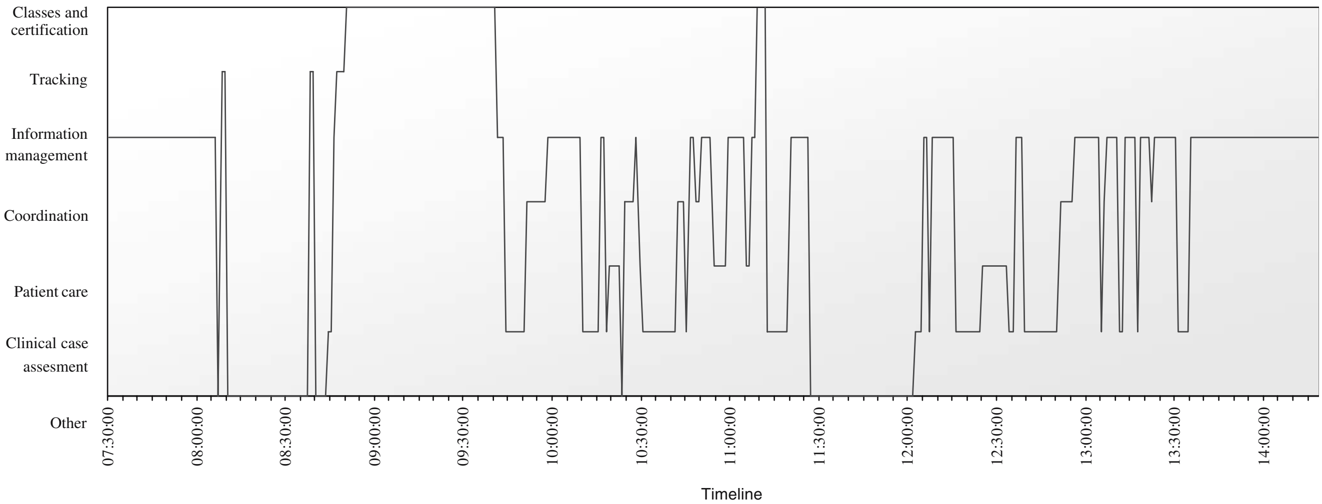

Figure 2 Switching between activities performed by a medical intern during a typical working day

by means of presentations (e.g., digital, paper-based), explain or discuss clinical cases and new techniques to improve the attendees' skills.

Due to the complex characteristics of hospital work, hospital staff needs to conduct activities that support others related to the clinical care of patients. For instance, hospital workers often have to engage in *information management* activities consisting on gathering and monitoring the status and availability of information required to conduct their work. Similarly, hospital workers regularly need to locate colleagues and medical equipment, which we refer as the activity of *tracking*. Another activity which plays a supporting role is *preparation*. This activity includes the actions performed to set up the environment before patient care is provided, such as the preparation of medicines or material.

The data collected was then transcribed into spreadsheets with columns for time stamps, description of actions, subjects involved in the actions, artifacts used, physical location, et cetera. We estimated the total time (per day per subject) that hospital workers spent on the activities established in the qualitative analysis. These results are presented in Table 1.

As described in Table 1, on average, most of the time is spent in information management (20.17%), followed by clinical case assessment (19.26%) and coordination (16.21%). An important part of the work of those observed is focused on activities that have a role that is secondary to patient care, such as information management, coordination, tracking, and, finally, preparing the patient census.

When analyzing the information per role, there are indeed differences regarding how much time each type of worker (nurses, physicians or interns) invests in performing these activities. We found that, the time nurses and medical interns spend performing administrative duties (e.g. information management, tracking, etcetera) surpasses the time they invest providing clinical care (e.g. patient care, clinical case assessment, et cetera). In contrast, physicians spend more time evaluating clinical cases and visiting patients. This happens because nurses and medical interns spend considerable time "*setting up*" the environment for physicians so they can focus on their main goal: providing clinical care. In addition, another significant difference is the nature of the activity being performed. For example, although, physicians, interns and nurses provide patient care, the actions executed by each role are quite different. Nurses are in charge of monitoring vital signs, administering medication and so on, while physicians and interns are in charge of providing specialized care such as catheter insertions.

Another characteristic we observed is that some of the activities performed by hospital workers are fragmented due to the need for tracking people or documents necessaries to accomplish their activities, taking care of emergencies or by interruptions. This fragmentation demands from hospital workers to continually switch between tasks regulating the time they last. As Fig. 2 shows, the most fragmented activities are coordination and tracking with 5 min of approximated duration; while, clinical case assessment, classes and certification and, preparation of medicines are conducted for longer periods of time (20–50 min approx.). Information management and patient care could be either fragmented or not depending on the location of the artifacts needed to perform those tasks. For example, while nurses are taking a patient's vital signs, they constantly need to move throughout the hospital premises looking for thermometers and other medical equipment, thus interrupting the activity being performed to gather equipment or information. Also, another additional feature shared by the activities performed by hospital workers is the location where they take place. We observed that these activities could take place either at a "base

Table 1 Time hospital workers spend performing different types of activities

Activity (time per day per subject)	Nurses			Medical interns			Attending physicians			Average total time	% of the entire day
	Mean time of segment	Average time	%	Mean time of segment	Average time	%	Mean time of segment	Average time	%		
Information management (IM)	0:03:42	1:18:08	18.15	0:07:19	1:50:54	27.82	0:03:29	0:58:26	16.62	1:22:29	20.86
Clinical case assessment (CCA)	0:01:52	0:22:18	5.18	0:03:35	1:06:04	16.57	0:03:14	2:07:54	36.39	1:12:05	19.38
Coordination (C)	0:02:11	0:41:18	9.59	0:03:29	1:09:09	17.35	0:02:35	1:19:42	22.68	1:03:23	16.54
Personal activities (PA)	0:04:49	1:03:49	14.82	0:04:14	1:20:24	20.17	0:03:10	0:42:30	12.09	1:02:14	15.69
Patient care (PC)	0:07:55	2:14:08	31.15	0:09:00	0:27:59	7.02	0:04:27	0:30:37	8.72	1:04:15	15.63
Preparation (PM)	0:02:37	1:19:42	18.51	0:01:14	0:03:28	0.87	0:00:27	0:02:12	0.62	0:28:27	6.67
Tracking (T)	0:00:58	0:05:30	1.28	0:02:11	0:18:17	4.58	0:01:18	0:09:59	2.84	0:11:15	2.90
Classes and certification (CC)	0:05:42	0:05:42	1.32	0:22:25	0:22:25	5.62	0:00:00	0:00:00	0.00	0:09:22	2.32
Unknown	0:00:00	0:00:00	0.00	0:00:00	0:00:00	0.00	0:00:07	0:00:07	0.03	0:00:02	0.01
All	0:29:46	7:10:35	100	0:53:27	6:38:40	100	0:18:47	5:51:27	100	6:33:34	100

location” (i.e., main place where subjects spend most of their time) or “on the move” (i.e., outside of their base location). For nurses, the base locations were either the nurse pavilion or the nurse office, whereas for medical interns and physicians it referred to a common office shared with other personnel of the department. Finally, the artifacts used and the people engaged in the execution of the activity are relevant.

2.3 Activity transitions

Table 2 shows the activity transition matrix. The rows in this matrix correspond to the activities executed by hospital workers and the columns correspond to the state transition (i.e., from one activity to another). The matrix was calculated per timestamps. A new timestamp was generated when new events altered one or more contextual variables. For example, if somebody arrives to the scene when a physician is discussing a clinical case with a colleague a new timestamp was generated. Those changes in the context include an individual’s location; artifacts or people the subject interacts with; or any action executed by the individual.

As the shaded values in Table 2 show, the activities present a recurrence phenomenon, which means, that the probability of remaining executing such activity, when the context changes, is higher than switching to execute another one. We found that the location where the activity is executed, and to a lesser degree its duration, are the most significant factors affecting activity recurrence. For exam-

Table 2 Percentage (per day per role) of the activity shift

Activities (% per day per informant	CCA	PC	C	P	IM	T	CC	PA
CCA	**50.55**	9.89	11.27	5.37	6.85	2.39	0.04	13.64
PC	21.23	**40.63**	11.93	6.06	5.97	2.54	0.00	11.63
C	18.64	8.32	**45.92**	5.77	11.26	4.91	0.18	5.00
PM	19.04	22.48	7.69	**39.61**	1.84	2.02	0.00	7.32
IM	18.99	8.03	15.30	6.16	**45.29**	4.63	0.05	1.55
T	14.53	12.80	15.41	7.88	11.75	**32.62**	0.00	5.00
CC	1.34	1.73	0.00	0.00	4.02	0.00	**89.36**	3.55
PA	14.20	11.84	4.73	7.10	1.42	4.73	3.31	**52.66**

ple, classes and certification is the activity with the highest level of recurrence. As Fig. 2 illustrates this activity was conducted from approximately 1 h and it is executed at base locations, such as meeting rooms or offices. In contrast, tracking is conducted by navigating the hospital premises for periods of time that average 51 s.

As we discussed, each activity performed by hospital workers is dependent on contextual variables such as the user's location and identity, the time of the day, the people with whom they collaborate, and the artifacts used during the execution of the activity, as well as, the activity duration. From this, we decided to design an approach to estimate the hospital worker's activities based on such contextual information.

3 Estimating user activities with neural networks

As mentioned before, activities are dependent on several contextual variables, which differ in the way in which they are codified for analysis. For example, location could be codified by assigning each place a number that represents it, whereas artifacts used were codified using a binary representation: one if it is being used and zero if not being used while performing the activity. Thus, inferring and managing these contextual variables is not trivial for a computer. Back propagation neural networks (BPNN) can manage several contextual variables, codified in different ways, as inputs or outputs, thus allowing to deal easily with this complexity. Hence, we decided to use them as the inferring engine to estimate the users' activity, by mapping from these contextual variables to activities.

3.1 Use of a supervised pattern recognition algorithm

Neural networks are models that have a significant performance on the recognition of complex patterns. A neural network is, more specifically, a supervised non-parametric model that learns from training examples, they can learn to map input sequences (artifacts being used, time, collaborators and location) to output sequences (activities). Once trained, the neural network can be used to classify incoming patterns into labeled classes [11].

Neural networks are basically formed by neurons, which are the processing units, and synaptic weights, which form the "memory" of the network. The neurons are organized in layers: one input layer, one or more hidden layers and one output layer. As it can be seen on Fig. 3, the units of each layer are interconnected with all of the neurons from the following layer using synaptic connections, which contain the synaptic weights. Also, every neuron has an activation function used to process data. As mentioned, the NN learns from examples: every time a training input pattern is presented to the network, an output pattern is computed, which is compared then to the correct output pattern. The difference (error) between the estimated and the correct patterns is then back propagated and the synaptic weights are adjusted. This method should be repeated for all the training samples until an acceptable error level is achieved.

We decided to use networks with 16 hidden units as this configuration provided a good compromise between accuracy and performance for a PDA. The learning method we used was backpropagation, which has several variants from which we utilized the Bayesian Regularization algorithm as it provided the best results. As activation function we used the sigmoid function on the hidden layers and the identity function on the output layer. The network was trained using the Matlab Neural Network Toolbox.

3.2 Architecture of the neural networks

Three neural networks were trained to estimate the activities, one per role (physicians, medical interns and nurses) since putting them all together added extra complexity and low performance when estimating the activity in the PDA (i.e., extra calculations should be carried out). Moreover, the activities performed by each role were clearly identified. For instance, they all perform the same activities (e.g., patient care), however they generally perform different medical procedures and use different artifacts.

As Fig. 3 shows we used five contextual variables to train the network. Four of them—location, artifacts, role and time—were used as inputs, whereas activity was used as output. The code for each contextual variable was taken from our qualitative analysis. As Fig. 4c illustrates, this information was reported in an observation format, in which, based on the qualitative information recorded (referred to as "what they were doing") each contextual variable was codified following our coding scheme. For example, activity was coded by assigning a representative tag to each activity identified (i.e., patient care, clinical case assessment). As Fig. 4c illustrates, the "IM" tag, in the last column of the first row in the observation format, indicates that the activity being executed is *information management*. As this row also shows, we can see that a nurse while performing such activity, at 8:57 she is in the hallway, using a nurse chart and office accessories.

Next, the information coded in the observation format is transformed into inputs and outputs to the corresponding neurons, as Fig. 4b illustrates. The network uses, as inputs, 1 neuron for time of the day, 1 neuron for user location, 1 neuron for every significant artifact the role might use and, 1 neuron for each role with whom they might collaborate; and, for outputs, 1 neuron for every possible activity performed by the role. For example, location was coded by assigning a representative number to each of the operation

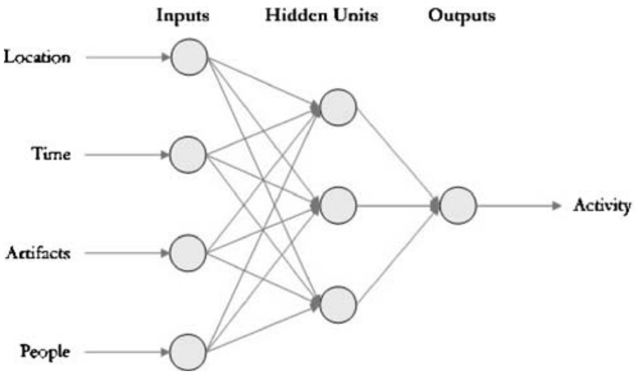

Figure 3 A simplified view of neural network used to estimate work activities

centers identified, such as bed wards, the hallway, personal area, offices and so on [23]. We identified 8 operation centers. Thus, in Fig. 4b the location neuron presents a number between 1 and 9 representing one of the operation centers where hospital workers' perform different activities. Whenever an artifact or collaborator is involved in certain activity the corresponding input neuron receives a value of 1 (one), otherwise, the input value is 0 (zero). For instance, as Fig. 4b illustrates, the values to the corresponding neurons for the inputs and outputs stand for the first row of the observation format. As the Figure shows, the time neuron have the time value on a decimal scale, the loc neuron have a 3 (i.e., hallway), the nobody neuron has the value of 1 which indicates that no interaction during this activity occur and; the AO and HE neurons have the value of 1 indicating that the artifacts being used are office accessories (i.e., AO) and the nurse chart (i.e., HE). And finally, the IM neuron has the value of 1 indicating that the activity being executed is information management (i.e., IM).

The number of neurons used for artifacts and people varies depending on the role. For example, to predict if a physician is formalizing notes, a relevant artifact is a computer or a typewriter. Nurses do not have access to such artifacts, thus there is no need to include a computer as a relevant artifact to train the nurse's network. Thus, the architectures used to train the networks are role dependant. For all three roles we used a single hidden layer with 16 neurons. Hence, the resulting networks are composed of

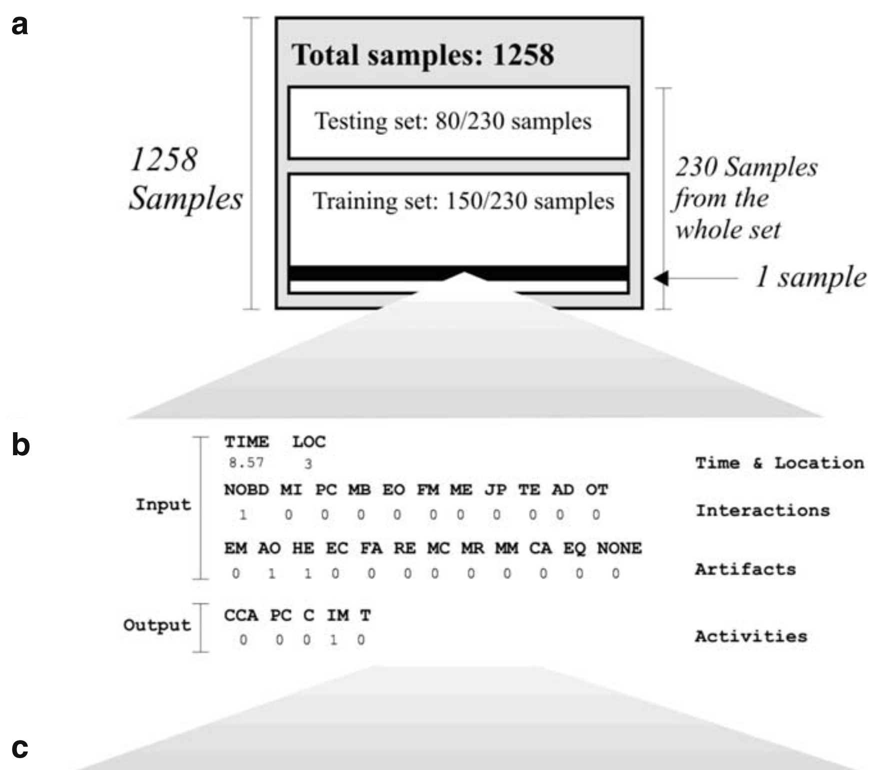

Time	What they were doing	Artifacts	Participants	Location	Primary Activity
08:57:30	She reports the patients' diagnosis in the nurse sheet	Nurse Sheet (NS)+ Office Accessories (OA)		Hallway	IM
07:36:01	She moves to the 240 room. She takes the patient's temperature and heart rate	Medical Material (MM)	Patient (P)	Bedwards	PC
12:21:00	She discusses the patient's diagnosis with the medical interns	Medical Record (MR)	Physician + Intern	Bedwards	CCA

Figure 4 **a** Sampling example from a physician for clinical case assessment. **b** Sample from the training set. **c** Coded fragment of the observation report three layers, with the following number of neurons: 22–16–7 for nurses, 27–16–6 for medical interns and, finally, 25–16–5 for attending physicians.

3.3 Codification and pre-processing of the data sets

As we discussed, the amount of time hospital workers spend in each activity is role dependant. Thus, for particular roles we could not gather enough samples to train the network. For example, for the activity of preparation we gathered 41 samples for interns and 36 for physicians, while for nurses there were 663 samples. Since this activity is similar to others such as patient care, having this small number of samples to train and test the network, influenced the way the network adjusted the synoptic weights, affecting thus, the memory of the network. In addition, the total percentage of the execution time of such activities (i.e., preparation, tracking and classes and certification) were minimal, therefore excluding such activities does not affect the way an application will work in a real working setting. Hence, we decided to eliminate from our experiments these activities. For all three roles the activity of tracking was eliminated. For the medical interns we also eliminated preparation, and for physicians we eliminated preparation and classes and certification. Table 3 shows the total samples obtained from the qualitative analysis (discussed in Section 2) and the samples used to train and test the network.

All these samples were grouped by activities which were divided into two data sets: one for training and another one for testing the network. To avoid a bias due to a larger amount of training data available for some of the more common activities such as patient care (PC), we decided to balance the training and testing sets. For instance, Fig. 4a shows how the total samples (1,258) were reduced to 230 samples for a physician performing clinical case assessment (CCA). Once the data sets were reduced we used 65% (150 samples) of the data for training and 35% (80 samples) for testing.

We conducted several experiments using different configurations of the network as well as a different size of the training set, however due to lack of space we present only some of them. In the following section we discuss our results.

4 Results and discussion

This section presents the results of our approach showing a confusion matrix per role, discussing the strategies to increase accuracy and the limitations of our approach.

4.1 Confusion matrices by role

We present our results in the form of confusion matrices. The rows in this matrix correspond to the values of the real activity performed (the activity codified in our observation format), and the columns correspond to the values estimated by the pattern recognition algorithm. The values in the diagonal are instances when the neural network correctly estimated the activity, whereas the rest represent the times when the neural network failed to guess the activity (i.e., confusing one with another). The average percent error is computed by summing up the non-diagonal values and the result is then divided by the total samples used to test the network.

Table 4 indicates the confusion matrix of the nurses' activities with a percent average error of 27.49%. The activity corresponding to classes and certification is the activity most accurately estimated (100% accuracy), followed by patient care with 77.50% and preparation with 75%.

Table 5 presents the confusion matrix of the medical interns' activities with an average percent error of 33%. As it shows, information management is the most accurate activity estimated (73.75%) followed by classes and certification (72.73%) and coordination (67.50%). Upfront, it is clear that the estimation accuracy for the medical interns is lower that obtained for the nurses.

Table 3 Number of samples used to train and test the neural network

A-ID	Nurses		Medical interns		Physicians	
	Total sample	Train/test	Total sample	Train/test	Total sample	Train/test
Clinical case assessment	608	150/80	479	150/80	1258	150/80
Patient care	899	150/80	156	90/60	287	150/80
Coordination	347	150/80	513	150/80	592	150/80
Preparation	663	150/80	36	–	41	–
Information management	121	72/40	667	150/80	339	150/80
Tracking	69	–	176	–	136	–
Classes and certification	28	15/10	36	15/10	–	–

Table 4 Confusion matrix for the estimation of activities performed by nurses

A-ID	CCA	PC	C	PM	IM	CC
CCA	**66.25**	10.00	13.75	3.75	6.25	0.00
PC	1.25	**77.50**	5.00	16.25	0.00	0.00
C	7.50	13.75	**73.75**	5.00	0.00	0.00
PM	7.50	2.50	13.75	**75.00**	1.25	0.00
IM	16.67	2.38	0.00	19.05	**61.90**	0.00
CC	0.00	0.00	0.00	0.00	0.00	**100**

Train 150, test 80, pct. Error. 27.49%

Table 6 shows the confusion matrix of the physicians' activities with an average percent error of 28.75%. As it shows, patient care is the activity most accurately estimated (78.75%), followed by coordination (72.50%) and information management (70%). Some activities, such as information management, are estimated with less accuracy, because some activities share several characteristics such as the place, time of day, people with whom they are conducted, and the artifacts used. For example, as illustrated in Table 4, our approach incorrectly estimated the activity of information management 16.67% of the times by confusing it with that of clinical case assessment. In addition, by analyzing the information per role, we can see differences in the level of accuracy with which some activities are estimated. For instance, for nurses, the activity of patient care was estimated correctly 77.50% of the time, for interns 67.27%, and for physicians 78.75%. In this case, we observed that the task switching experienced by each role affects the way in which this activity is estimated. For example, as Table 4 shows, the activity of patient care was taken for coordination 13.75% of the times. We observed that, for a short period of time, medical interns and physicians constantly switch between these activities. In these cases, the network does not have enough evidence to differentiate between them, since the contextual information (of the user and the environment) does not change during such period. However, we might be able to differentiate between them by identifying particular attributes of each activity (such as its duration).

4.2 Strategies used to increase accuracy

We applied different strategies to increase the estimation accuracy. We observed that some of the estimated activities were selected with some degree of uncertainty. By comparing those Uncertain Estimated Activities (UEA) with the real activities executed, we found that some of activities are performed by extended periods of time presenting a recurrence phenomenon, this is that the probability of staying executing an activity is higher than changing to another one. For example, a physician assesses a clinical case for an approximate duration of 15 to 20 min, hence it is very likely for him to keep on this activity rather than switching to information management, Thus, if the neural network estimates that after a clinical case assessment the physician changes activity to engage in information management, this estimation might be inaccurate. This could be detected by comparing the current values for each of the estimations. If the one corresponding to information management has a higher value but still close to clinical case assessment, it might be safer to assume that the activity has not changed. By taking advantage of the UEA duration, we have a greater probability that the hospital workers will keep performing those UEA rather than changing to another activity, we call this the *recurrence principle*. To cope with this, we first identified those UEA and then, by taking into account the UEA duration we applied our recurrence principle.

The approach we used to identify the activity estimated is to select the output unit of highest value. Despite this, sometimes selecting the highest value does not provide enough evidence to assure that the selected activity is the real one, resulting thus, in a UEA. To identify the UEA two strategies were used. The first strategy involves the identification of those activities corresponding to neurons which activation values are less than 0.55. In this case, there is no strong evidence supporting one activity (i.e., a dominant output is missing). In the same way, confusion or

Table 5 Confusion matrix for the estimation of activities performed by medical interns

A-ID	CCA	PC	C	IM	CC
CCA	**58.75**	12.50	15.00	8.75	0.00
PC	9.09	**67.27**	18.18	3.64	0.00
C	12.50	5.00	**67.50**	15.00	0.00
IM	7.50	2.50	15.00	**73.75**	0.00
CC	9.09	0.00	18.18	0.00	**72.73**

Train 150, test 80, pct. Error. 33%

Table 6 Confusion matrix for the estimation of activities performed by physicians

A-ID	CCA	PC	C	IM
CCA	**63.75**	10.00	12.50	13.75
PC	2.50	**78.75**	13.75	5.00
C	11.25	5.00	**72.50**	11.25
IM	11.25	7.25	12.50	**70.00**

Train 150, test 80, pct. Error. 28.75

uncertainty exists when two neurons have similar activation values that are close to 1 (e.g., two dominant activities with values 0.70 and 0.67). In this case, we calculate the difference between the values of the activity selected and each of the outputs rejected. If such difference is less that 0.12 we consider this activity estimated as an UEA.

Once the UEAs are identified, we selected those activities which present the recurrence phenomena by calculating each activity average time of duration. Although, clinical case assessment and classes and coordination present this recurrence because of their nature, each activity level of recurrence depends on the role of the user who executes it. For example, although information management, for interns and physicians, is performed for longer periods of time, nurses generally switch between patient care and this task, reducing its average duration. For instance, while a nurse is taking the vital signs of a patient, she reports in the patient chart his temperature, pressure and heart rate. Thus, the nurse conducts both activities by frequently switching between them. The activities which present the recurrence phenomena, in addition of clinical case assessment and classes and certification, for nurses are coordination and preparation of medicines and material and; information management in the case of medical interns. Once the UEAs were identified and the recurrence principle was enforced we reduced the estimation error by approximately 5%.

4.3 Limitations of the approach

The results obtained are constrained by the data used to train the neural network. Although, the data is exhaustive, it was collected in the internal medicine area of a mid-size hospital. Hence, it is not possible to derive findings that can be directly extended to other hospitals. However, the activities performed by hospital staff are in general the same as those in other hospitals in Mexico and other countries as reported in [4, 31], thus we believe that the general findings that we obtained are not restricted to this hospital in particular.

The technique used to gather the data requires extensive field work. This is a disadvantage of the approach. However, a smaller field study could be made to tune the classifier to a new setting. In addition, we believe that our results could lead the way to refined techniques that could reduce the effort made for data gathering, such as semi-automatic data collection. In this regard, to facilitate the collection of samples at the field site an application for a PDA could be used. This application could have a map of the area, icons that represent roles and artifacts, and a simple user interface where a researcher can tap on a specific location and a set of icons while selecting, from a previously defined activity list, the activity being executed. This will reduce the time spent by researchers transcribing information and formatting data to generate statistics.

Moreover, implementing our solution would require the deployment of technology within the hospital to infer the location of people and the artifacts being used to estimate the activity of users. However, some of these technologies are gradually being introduced in hospital environments. Once the contextual data is available, the activity estimation can be done in real time with the approach proposed. In fact, the challenge might be to devise an appropriate interface, not to bother the user with very sudden changes. The time required to estimate the contextual variables will depend on the specific technology being used. One of the more challenging contextual variables to estimate, location, can be estimated quite rapidly through several techniques [8, 17, 20].

Finally, we estimate the activity by taking into account contextual information, such as location and artifacts being used, which would be estimated through other techniques or by using sensors. These methods are also prone to errors which would affect the accuracy of the estimation. However, most of the errors incurred in the estimation of the contextual variables will not have a significant impact on activity estimation. For instance, the distance between operation centers is normally much larger than the average error of current commercially available location estimation approaches.

5 Comparing the neural network with the activity estimation made by expert observers

In order to evaluate the difficulty of estimating hospital workers' activities and establish an important basis for comparing our approach, we conducted an experiment to determine how people, familiar with hospital work, estimate activities using contextual information. We provided six users with information of the contextual variables (role, location, time, artifacts used and presence of colleagues) used by the neuronal network to infer activities. And then, evaluated the precision with which such users were able to infer those activities. Six researchers participated in this experiment, including one of the authors. Those researchers were the ones who followed the hospital workers during the case study—completing 80 h of observation (on average). Thus, they were familiar with hospital work and the coding scheme used. To form our *human test set*, we took for each activity, an amount of samples proportional to the ones used to test the neuronal network, using a total 120 samples (40 per role). This set of samples was presented in a form similar to the observation report, that we call the *testing report*. We asked each researcher to situate herself in the description shown by the testing report and based on the evidence infer the activity being performed. For instance, a sample of the testing report might illustrate that the researcher: (1) is an attending physician, (2) that is noon, (3) that she is in the bed wards, (4) that she is with medical interns and (5) that she is using a medical record and a nurse chart. From five possible answers—that correspond to the activities such as clinical case assessment, patient care, coordination and so on—the researcher must infer which activity is being performed.

Table 7 shows the confusion matrix of the nurses' activities with an average error of 48.05%. As Table 7 shows, the activity corresponding to information management is the most accurately estimated (86.67%), followed by patient care (66.67%) whereas clinical case assessment was the activity less accurately estimated (15%). It is interesting to note that the mistakes made by the researchers are not random. Rather, the mistakes appear to include a certain amount of bias related to the evidence used to infer hospital workers' activities. For example, if a user is using office accessories it is quite obvious that the activity being performed would be information management, or if a user is using medical equipment the activity performed would be patient care. Hence, some contextual variables are highly dominant.

The expert observers could correctly estimate the activities 48.82% of the time (on average). They were most accurate in estimating nurses' activities (51.95%), followed by physicians (47%) and medical interns (48%). This could be partially explained because the differences between the evidence used (such as the people involved and the artifacts used) to estimate each activity is role dependant. For example, for nurses, the tasks related to patient care frequently involved the use of medical equipment, the presence of the patient to take his vital signs, while for physicians and interns those tasks often involved the use of information artifacts such as the medical record to explain to a patient his condition. However, physicians and interns might also use information artifacts in collaboration with the patient to assess his condition. Hence, researchers in the role of physicians or interns could easily confuse clinical case assessment with patient care. Upfront, it is clear that the precision of these results is significantly lower than the ones obtained with the neural network.

6 Activity-aware applications

Although, several efforts have been made to estimate activities with a low level of abstraction such as if a user is walking, sitting [22] or chewing [5]; this information is not enough to inform a context-aware application how to adapt its behavior to support the work performed in hospitals. Thus, we need to estimate activities with a higher level of abstraction such as, if a physician is evaluating a clinical case or providing clinical care. By knowing these activities we can discover the contextual information relevant to the task at a hand or infer secondary context such as user availability.

In this section we discuss how the results obtained with the activity estimation method proposed can be applied in activity-aware applications. In particular, we discuss how our results could be applied to enhance the interface of a voice-based hospital communication system and providing quicker access to relevant hospital information.

6.1 Activity-aware collaboration

Mobility and collaboration generate a need to contact colleagues within the hospital, either to discuss a case with a specialist or request help to transfer a patient. Several mechanisms are used for these purposes and technology has been developed to assist in this task; such as the Vocera communication system which enables users to contact a fellow hospital worker either by name, role or location using a hands-free voice communication system. The problem is that these systems are largely unaware of the social situations surrounding their usage and the impact that their actions

Table 7 Confusion matrix of estimations made by expert observers on the activities performed by nurses

A-ID	CCA	PC	C	PM	IM	CC
CCA	**15.00**	27.50	20.00	5.00	32.50	0.00
PC	2.50	**67.50**	10.00	17.50	2.50	0.00
C	20.00	25.00	**50.00**	0.00	2.50	2.50
PM	2.50	12.50	17.50	**62.50**	5.00	0.00
IM	16.00	0.00	0.00	0.00	**84.00**	0.00
CC	6.67	13.33	40.00	13.33	0.00	**26.67**

Test 40, pct. Error. 48.05%

have on these situations. If the system could be aware of the user's availability, they could use this information to negotiate interruptions at appropriate times, improving thus, human computer interaction [14]. Availability is information that can be derived from knowing the activity being performed by a person. For example, when hospital workers' are involved in clinical case assessment, patient care, classes and certification and even, preparation they, in general, do not want to be interrupted. We recurrently observed that medical workers, especially interns, wait until the discussion finishes in order to approach a physician. In contrast, when they are engaged in other activities, such as information management or coordination, the level of their interruptibility is higher. Hence, we decided to group the activities based on the availability perceived by hospital workers. The non-available activities correspond to the activities with the aim of improving the quality of patient care. Those activities include clinical case assessment, patient care, classes and certification and preparation. While the others activities such as information management and coordination are the ones with a higher level of interruptibility.

In Table 8 we present the activity estimation results obtained with our approach, but grouped by those activities that can be associated with the user's availability. The table shows that the estimation accuracy ranges from 70 to 90%. A system such as Vocera could use these results to decide whom to interrupt and when. For instance, if a physician needs the help of a nearby nurse, he can make this request and the system will decide, based on the availability estimation, which nurse to call at that particular moment.

6.2 Activity-aware information retrieval

Hospital Information Systems manage large amounts of information, both clinical and administrative. This includes patient records with laboratory results and medical images; medical guidelines and procedures; staff assignments; reference material, etc. Navigating through this information can be time consuming when the users want to consult a single piece of data. It has been suggested that a context-aware application could display the medical record when a physician is in front of a patient's bed [25]. However, the physician could be there preparing medical equipment to perform a catheter insertion or prescribing medicine. In this case, he might want to consult medical guides or pharmacological databases instead of the patient's health record. Thus, location information might not be enough to decide which information to present to the user or make more available. It is hard to stipulate what contextual information is relevant and how to adapt context-aware applications to present services and information that are suitable to the users' current activity [16].

The users' activity can be used to identify the information relevant to the action at hand and differentiate between useful and unwanted data (with many degrees in between) and use this differentiation to meet goals, such as timely care delivery. Therefore, to identify when, depending on the activity being performed, is more relevant to consult medical information (e.g., medical record, laboratory results) than other information (e.g., the location of colleagues, journal articles), we classified the activities for which medical information is needed. For example, when hospital workers' are assessing a clinical case, caring for a patient or managing information they often consult that patient's medical record. However, when they are coordinating, tracking or in classes they do not use such information. Table 9 shows the estimation of the need to consult patient medical information based on the hospital worker's activity. As in the case of hospital worker's availability the estimation accuracy increases when the activities are grouped according to this parameter.

Figure 5 shows the interface of a Hospital Information System. The upper frame shows the application with which the user interacts. Below, the application displays two thumbnails of screens suggested to the user based on his current activity. If the user selects one of the thumbnails, the corresponding application will appear in the main screen.

Table 8 Confusion matrix for estimating the availability of hospital workers

	Nurses		Interns		Physicians	
	NA	A	NA	A	NA	A
NA	**89.58**	10.41	**76.59**	23.40	**90**	10
A	20.61	**79.38**	13.83	**86.16**	29.11	**70.88**

Error: nurses 14.01%, interns 18.33%, physicians 14.73%

Table 9 Confusion matrix for the estimation of hospital workers' need for patient medical information

	Nurses		Interns		Physicians	
	NR	R	NR	R	NR	R
NR	**90.78**	9.21	**83.73**	16.26	**87.02**	12.97
R	23.59	**76.40**	29.67	**70.32**	27.5	**72.5**

Error: nurses 12.66%, interns 20.33%, physicians 16.61%

6.3 Activity-aware task switching

The mobileSJ system was implemented to assist mobile users in the management of their multiple activities and collaborations [7]. This application implements the concept of a working sphere [15] in its computational representation: an e-sphere. A working sphere is a concept introduced to conceive the way in which people organize and execute their work activities. mobileSJ allows the user to manage their multiple activities and their information and contextual resources while away from the desktop, since it runs on a PDAs or SmartPhone. The user can define e-spheres for each of his activities and associate to them, information resources, contacts relevant to the activity, emails related to the activity and pending issues (Fig. 6). When a user switches between e-spheres, each e-sphere is enabled to quickly gather and retrieve its own workspace state (windows positions,status and overlay order) and context information like opened documents, idle time, et cetera in a silent manner. In addition, mobileSJ allows sharing activities and resources as well as communicating with colleagues through either SMS messages or phone calls.

Given that hospital workers need to cope with multiple activities, which are often fragmented by interruptions, and which require them to gather and consult a variety of information resources, mobileSJ can provide users with mechanisms to easily manage their activities and their associated resources helping to achieve the necessary context when switching between activities. However, users must identify their working spheres and explicitly specify the resources (i.e., contacts) associated to that sphere, in order to obtain the benefits provided by mobileSJ. Also, a user must explicitly select the active sphere relevant to the task at a hand. By knowing the user's current activity mobileSJ can proactively identify the relevant sphere, at a given moment. Thus, by taking into account the information of the activity being performed by hospital workers, mobileSJ could retrieve the e-sphere relevant to the task at hand. For example, a physician might encounter a colleague in the hallway, where they decide to discuss the clinical case of a patient both have in common. In this case, neither the hour of the day nor the hospital workers' contextual information (e.g., identity, location) could inform which sphere to retrieve; moreover, what information associated to such sphere is relevant. However, by knowing that both hospital workers are discussing a clinical case we can infer that they must want to share the patient's laboratory results to enrich the discussion. In this case, mobileSJ will retrieve the relevant e-sphere and put the associated resources within easy reach of the users. Similarly, the application could suggest the users to append newly created resources to the corresponding e-sphere.

Figure 6 Activity-aware information access to the patient's medical record

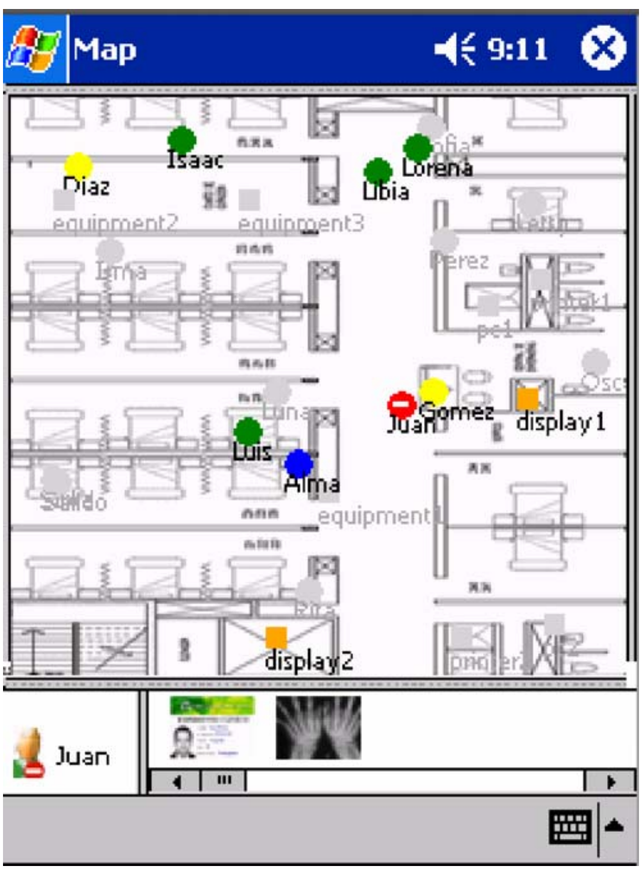

Figure 5 Activity-aware information access to the patient's medical record

7 Sensing contextual information required for activity estimation: issues and opportunities

One of the biggest challenges for the deployment of activity-aware applications is the availability of robust methods to estimate hospital work activities along with the technology necessary for such estimation. In our approach four contextual variables must be measured: location, time, artifacts used and the people with whom users collaborate. Several methods and technologies have been proposed to capture and monitor this contextual information ranging from the use of simple sensors such as RFID tags to complex systems that fusion the information from different sources. However, due to the complex characteristics of the work in settings such as hospitals, new issues need to be addressed for the efficient operation of these technologies in real applications. In this section, we discuss how technological approaches proposed in the pervasive computing literature provide evidence of the feasibility of the approach proposed here, as well as, revealing new unforeseen opportunities for this field.

In regards to estimating the location of users, considerable research has been conducted in the past few years. Several approaches, techniques and sensing technologies have been used such as ultrasonic pulses, infrared, radio-frequency signals, high-precision cameras, pressure sensors, and RFID. As a result, there is an ample range of solutions that provide to some extent reliable estimations ranging most of them from 1–2 m [17], which might be sufficient for many applications. However, along with the recent boom of sensor networks the problem of locating nodes has become of particular interest as sensor networks could eventually provide enhanced solutions to the location problem on topics such as scalability, low-cost, and decentralized and automated estimation [28]. For sensor networks, in particular, the information gathered by a node must provide location information in order to make sense, of course, depending on the nature of the application. Trends such as sensor networks, RFID along with the recently revealed Memory Spot Chip[1] by the HP Labs are opening the door for the deployment of fully pervasive applications. The latter one, in particular, provides many advantages such as its size (less than half the size of a grain of rice), 10 Mbps data transfer rates, no need for a battery and, finally, storage capacity of up to 4 Mb.

Despite that several approaches have been proposed to infer user proximity based on location sensing, proximity measurement, and device discovery [17, 19, 20, 26]; knowing that a user is close to another one does not provide enough evidence of actual collaboration. A step in the direction of estimating whether users are collaborating or not is to create ad-hoc networks where each individual node in the network can monitor its local region and communicate with other nodes to collaboratively produce a high-level representation of the overall environment. These ad-hoc network technologies enable devices to detect and connect to other devices that are in sufficient proximity, and form mobile peer-to-peer (P2P) networks. P2P network technologies are ideal for transmitting information and establishing application sharing between devices that are physically collocated [12]. Once an application has strong evidence that the user's devices are executing actions that involve collaboration (such as sharing the screen of a PDA), collaboration could be easily inferred. However, following an approach such as this one requires a priori calibration and configuration. In environments where rapid response to emergencies is required, approaches that work "out of the box" by detecting collaboration directly from sensors would be the adequate solution. For example, if a sensor could provide primary evidence that a conversation, a shared application session or an information transmission is taking place.

Several efforts have been made to detect the artifacts being manipulated by a user. These approaches range from attempts that require the user to wear complex equipment to methods that combine the advantage of easy sensor deployment with that of unobtrusive sensor detection. For instance, the RFID-detecting glove houses in a small box a SkyeTek RFID reader which communicates to an antenna the information captured [30]. Hence, the glove acts as an information proxy between the user and the artifacts with RFID tags attached. Thus, each time a user grabs a tagged artifact the glove will capture such information identifying the artifact being used at that particular moment. Another approach in this direction are sensors designed as "tape on and forget" devices that can be quickly and ubiquitously installed in real settings. In this approach, "stick-on" sensors with an accelerometer, clock, and local memory are placed on the objects [24]. Accelerometers detect touching, and record that to local memory. When the sensors are later removed, the data can be analyzed to reconstruct the set of all object touches. Alternatively, the two types of sensors describe above can be combined, by using the accelerometer from the first type of sensor to detect interactions and the RF protocol of the second one, to transmit this information wirelessly. This method has potential, but has the disadvantage of being incompatible with the commercial tags. Hence, techniques that operate by interacting with long-range readers have been proposed [13]. However, knowing that a user is holding a medical record does not provide enough evidence to judge if a user is engaged in information management or clinical case assessment. By knowing that the hospital worker uses a typewriter along with the medical record, our approach could infer that such user is performing activities related to information management. Hence, combining measures from multiple sensors could increase the confidence value for a particular interpretation. This approach requires saturating the hospital premises with pervasive sensors and technology that must assemble and transmit contextual information from different sources in the environment, resulting thus, in a wireless sensor network. This sensor network must have capabilities for context fusion providing reliable ubiquitous context by combining services in parallel (to offset noise in the signal) and sequentially (to provide greater coverage) [2].

[1] Memory Spot by HP Labs, http://www.hpl.hp.com/news/2006/jul-sept/memoryspot.html

Despite the advance in sensor networks, there are some issues that must the addressed before any wide-ranging implementation, such as cost and energy consumption as devices that are GPS and even Wi-Fi enabled still perform very poorly. The viability of projects based on a vast number of sensors will depend significantly on the cost of the devices (considerably low, cents perhaps), the replacement of sensors should be at least from 6–10 years, and the configuration of the network itself must be realized without any significant human effort [28].

8 Related work

Previous efforts have been made to infer the activities of people using different techniques. For instance, [22] (19) reports a dead-reckoning method to recognize and classify user's sitting, standing and walking behaviors by obtaining the acceleration and angular velocity of wearable sensors being bear by users. Activities have also been inferred by detecting the interaction of users with particular objects [30]. This is done by tagging every item of interest using RFID tags and reading them through a RFID-detecting glove. Similarly, as part of the MIT project house_n [18], the deployment of MITes (MIT Environmental Sensors) is used to monitor people's interaction with objects in the environment. All of these projects have a similar goal: the use of technology to support the Activities of Daily Living (ADL) and Instrumental ADL (IADL) of elders. In contrast, our work focuses on supporting the daily-life work activities of hospital staff which usually involve other sort of contextual variables completely different from those coming from the home setting.

As part of Georgia Tech AwareHome project, sound recordings were used to monitor and infer activities in the home environment [5]. Arrays of microphones were placed to infer activities such as meal preparation. Although sound could be used to detect the presence of a colleague, or the use of a particular artifact in a hospital, it will be difficult to directly infer activity from such information, given the complexity of he activities involved and the difficulty of controlling noise in such environment. In contrast with the very specific activities detected by this work (e.g., chewing), our aim is to estimate work activities at a higher level of abstraction.

The use of production rules has also been proposed to identify the activities performed by hospital workers [10]. The presence of artifacts and people with RFID tags attached to them are used to trigger the inference rules coded by a knowledge engineer. However, the creation of these rules is a time consuming task that requires considerable expertise and knowledge of the setting. Furthermore, consistency among these rules is not easily achieved. No results are reported on the accuracy of this method.

As mentioned in Section 5, from the activity being performed by a person one can infer its availability. With the use of sensors and considerations of social behavior the availability of an individual in an office environment as been estimated with an 80% accuracy [14]. The results are similar to the ones we obtained for a hospital. However, rather than inferring availability directly, we first determine the activity being performed, information that could also be used by an activity-aware application.

9 Conclusions

Estimating user activity, as we discussed, is a complex task and, despite the importance of knowing the user's activity has been highlighted in pervasive computing [1, 34], few attempts have been conducted to address this problem. In this paper, we presented an approach that uses neural networks to estimate hospital worker's activities. To train the network, we used the information recorded from a workplace study conducted in a hospital. By following this approach we could correctly estimate hospital workers' activities 75% of the time (on average). In addition, we discuss how once an application has strong evidence of the users' activity, it could adapt itself by displaying information relevant to the task at hand, and infer secondary context, such as availability. To illustrate this, we discuss how these results can be used in the design of activity-aware applications.

One of the biggest challenges for the deployment of activity-aware applications is the availability of robust methods to estimate hospital work activities; the latter along with the necessary technology to sense the contextual information with which such estimation is made. We discussed how technological approaches proposed in the pervasive literature provide evidence of the feasibility of our approach as well as revealed opportunities for the deployment of pervasive sensors. However, the characteristics of the work in real settings, such as hospitals, generate new challenges that need to be addressed for the efficient operation of this technology in real settings. These challenges include the improvement of the capabilities of sensors (such as maximizing their battery, transfer rate and storage), as well as, providing sensor and ad-hoc networks with capabilities to fusion context by combining measures from multiple sources to increase the confidence value for a particular interpretation.

We plan to explore other estimation techniques such as the Hidden Markov Models to improve our results. Our hypothesis here is that by taking into account information from the past, we can determine the probability of switching from one state (activity) to another. In addition, we want to use our results on the implementation of an activity-aware hospital application and evaluate its practical implications. This requires to take into account that activity estimation is often unreliable, especially when based on contextual information that can itself be uncertain, which raises the need for appropriate mechanisms for the identification, representation and management of uncertain information. We are currently working on an uncertainty management technique to deal with this issue. Such approach would prompt the user for feedback when uncertainty is beyond a certain threshold. Finally, the approach could be enhanced by incorporating information (i.e., estimations or sensing) from other nodes in the vicinity.

Acknowledgments We thank the personnel at IMSS General Hospital in Ensenada, México. This work was partially financed by CONACYT under contracts C03-42447 and U-40799, and through scholarships provided to Monica Tentori, Luis A. Castro and Elisa B. Moran.

References

1. Abowd GD, Mynatt ED (2000) Charting past, present, and future research in ubiquitous computing. ACM Trans Comput-Hum Interact (TOCHI) 7(1):29–58
2. Abowd G, Mynatt E (2002) The human experience. IEEE Pervasive Computing 1(1):48–57
3. Bardram JE (2004) Applications of context-aware computing in hospital work: examples and design principles. In: ACM symposium on applied computing. ACM Press, Nicosia, Cyprus, pp 1574–1579
4. Bardram JE, Bossen C (2005) Mobility work: the spatial dimension of collaboration at a hospital. In: ECSCW. Kluwer Academic Publishers, Paris, France, pp 176–185
5. Bian X, Abowd, GD, Rehg JM (2004) Using sound source localization to monitor and infer activities in the Home. Georgia Institute of Technology, Atlanta, GA, Technical Report; GIT-GVU-04-20
6. Bossen C (2002) The parameters of common information spaces: the heterogeneity of cooperative work at a hospital ward. In: CSCW. Kluwer, New Orleans, Louisiana, pp 176–185
7. Camacho J, Favela J, Gonzalez VM (2006) Supporting the management of multiple activities in mobile collaborative working environments. In: CRIWG. Springer, Valladolid, Spain, pp 381–388
8. Chen M, Haehnel D, Hightower J, Sohn T, LaMarca A, Smith I, Chmelev D, Hughes J, Potter F (2006) Practical metropolitan-scale positioning for GSM Phones. In: 8th international conference on ubiquitous computing. Springer, Orange County, CA, USA, pp 225–242
9. Chin T (2005) Untapped power: a physician's handheld. AMNews, Danville, KY
10. Christensen HB (2002) Using logic programming to detect activities in pervasive healthcare. In: International conference on logic programming. Copenhagen, Denmark, pp 21–436
11. Duda RO, Hart PE, Stork DG (2000) Pattern classification. Wiley, NY
12. Favela J, Tentori M, Markarian A, Castro LA, Amaya I (2006) Integrating heterogeneous devices in support of local mobility. In: 12th Americas Conference on Information Systems. Association for Information Systems, Acapulco, Mexico
13. Fishkin KP, Jiang B, Philipose M, Roy S (2004) I sense a disturbance in the force: unobtrusive detection of interactions with RFID-tagged objects. Ubicomp, Innsbruck, Austria
14. Fogarty J, Hudson SE, Atkeson CG, Avrahami D, Forlizzi J, Kiesler S, Lee JC, Yang J (2005) Predicting human interruptibility with sensors interruptibility with sensors. ACM Trans Comput-Hum Interact 12(1):119–146
15. González V, Mark G (2004) Constant, constant, multi-tasking craziness: managing multiple working spheres. In: CHI. ACM Press, Vienna, Austria, pp 113–120
16. Greenberg S (2001) Context as a dynamic construct. Hum-Comput Interact 16(2–4):257–268
17. Hightower J, Borriello G (2001) Location systems for ubiquitous computing. IEEE Comput:57–66
18. Intille S, Larson K, Kukla C (2000) House_n: The MIT home of the future project. Department of Architecture, Massachusetts Institute of Technology, Cambridge, MA
19. Krumm J, Hinckley K (2004) The NearMe wireless proximity server. Ubicomp, Innsbruck, Austria
20. Krumm J, Horvitz E (2004) LOCADIO: inferring motion and location from Wi-Fi signal strengths. In: First annual international conference on mobile and ubiquitous systems: networking and services. IEEE Press, Boston, MA, USA, pp 4–13
21. Lapinsky J, Weshler S, Mehta M, Varkul D, Hallett T, Stewart T (2001) Handheld computers in critical care. Crit Care 5(4): 227–231
22. Lee SW, Mase K (2002) Activity and location recognition using wearable sensors. IEEE Pervasive Computing 1(3):24–32
23. Moran EB, Tentori M, González VM, Martinez-Garcia AI, Favela J (2006) Mobility in hospital work: towards a pervasive computing hospital environment. Int J Electron Healthcare 3 (1):72–89
24. Munguia ET, Intille SS, Larson K (2004) Activity recognition in the home using simple and ubiquitous sensors. In: Pervasive. LNCS Springer, Heidelberg, pp 158–175
25. Munoz M, Rodriguez MD, Favela J, Martinez-Garcia AI, Gonzalez VM (2003) Context-aware mobile communication in hospitals. IEEE Comput 36(9):38–46
26. Naya F, Noma H, Ohmura R, Kogure K (2005) Bluetooth-based indoor proximity sensing for nursing context awareness. In: IEEE International Symposium on Wearable Computers, Washington DC, USA
27. O'Connor MC (2006) Testing ultrasound to track, monitor patients. RFID J 1(31):2
28. Patwari N, Ash JN, Kyperountas S, Hero Lii AO, Moses RL, Correal NS (2005) Locating the nodes: cooperative localization in wireless sensor networks. IEEE Signal Process Mag 22:54–69
29. Pentland A (2004) Healthwear: medical technology becomes wearable. IEE Comput 37(5):42–49
30. Philipose M, Fishkin KP, Perkowi M (2004) Inferring activities from interactions with objects. IEEE Pervasive Computing 3 (4):50–57
31. Reddy M, Dourish P (2002) A finger on the pulse: temporal rhythms and information seeking in medical work. In: Computer supported cooperative work. ACM, New Orleans, Louisiana, pp 344–353
32. Stanford V (2003) Beam me up, Doctor McCoy. IEEE Pervasive Computing 2(3):13–18
33. Strauss A, Corbin J (1998) Basics of qualitative research: techniques and procedures for developing grounded theory, ed. Sage, Thousand Oaks, CA
34. Weiser M (1998) The future of ubiquitous computing on campus. Commun ACM 41(1):41–42

Jesus Favela is a Professor of Computer Science at CICESE, where he leads the Collaborative Systems Laboratory and heads the Department of Computer Science. His research interests include ubiquitous computing, medical informatics, and CSCW. He holds a B.Sc. from the National Autonomous University of Mexico (UNAM) and M.Sc. and Ph.D. from the Massachusetts Institute of Technology (MIT). He is a member of the ACM and the American Medical Informatics Association (AMIA), and Former President of the Sociedad Mexicana de Ciencia de la Computacion (SMCC).

Víctor M. González is a Lecturer in Human-Computer Interaction at the School of Informatics of the University of Manchester, United Kingdom. He conducts investigations on technology usage in home, office and hospital settings and specializes on the use of ethnographic methods (e.g. interviews, participant observation) and participatory design techniques (e.g. scenario-based design). He received a Ph.D. and Master degrees in Information and Computer Science from the University of California at Irvine and a Master degree in Telecommunications and Information Systems from the University of Essex, United Kingdom. He is a member of ACM SIGCHI and Vice-President of the SIGCHI Mexican Chapter.

Monica Tentori is a Ph.D. student in Computer Science at CICESE. Her research interests include ubiquitous computing, HCI and medical informatics. She holds a M.Sc. in Computer Science from CICESE and a B.Sc. in Computer Science from the Autonomous University of Baja California (UABC) in Mexico. She is a student member of the SIGCHI special interesting group.

Elisa B. Morán works at the Telematics Department at the CICESE, where her focus is on the analysis and development of Information Systems. She holds a M.Sc. in Computer Science from CICESE and a B.Sc. in Computer Science from the Autonomous University of Baja California (UABC) in Mexico.

Luis A. Castro is a Ph.D. student at the School of Informatics of the University of Manchester. His research interests include ubiquitous computing, location aware computing, neural networks and HCI. Castro received a bachelor's in Information Technology from the Institute of Technology in Mexicali, Mexico and a master's in Computer Science from CICESE. He is a student member of the ACM.

Ana I. Martínez-Garcia is an Associate Professor of Computer Science at CICESE, where she studies the application of technology to organizational process reengineering specially in the health sector; she works on workflow, decision making, knowledge management and technology adoption. She received a Ph.D. and an M.Sc. in Advanced Computer Science from the University of Manchester, UK, and a B.Sc. in Computer Science from UABC.