

Case-based Reasoning for Situation-aware Ambient Intelligence: A Hospital Ward Evaluation Study

Anders Kofod-Petersen¹ and Agnar Aamodt²

¹ SINTEF ICT,
S. P. Andersens vei 15 b,
7465, Trondheim, Norway
akof@sintef.no

² Department of Computer and Information Science,
Norwegian University of Science and Technology,
7491 Trondheim, Norway
agnar@idi.ntnu.no

Abstract. Ambient intelligent systems are defined as being able to perceive their environment, being aware of the presence of people and other agents, and respond intelligently to these agents' needs. Today the hardware requirements for achieving these capabilities are met. Earlier work have argued that knowledge intensive case-based reasoning is a feasible method for ambient intelligence. In this paper that argument is supported by testing of an implementation in a hospital ward domain, which shows that despite some issues related to the current implementation the case-based reasoner performs at an acceptable level.

1 Introduction

As computers are becoming more *ubiquitous* [1], *pervasive* [2], and *ambient*¹ [4], the need for methods that assist their users in smart and intelligent ways is rapidly increasing. In an ambient system an actual physical computer device will be situated in various physical and social environments at different times, and a challenge is therefore to identify and interpret that environment. Ambient *intelligent* systems, as defined by the ISTAG group [5], are characterised by being able to perceive their environments, be aware of the presence of people and other agents, interpret their own role in that context, and respond intelligently to one or more agents' needs [6].

Realising ambient systems relies on miniaturisation of technology, high calculation power and interconnectivity. According to Satyanarayanan [7], all these hardware tools are currently available. Yet, many of the scenarios described in

¹ Recently the term *everywhere computing* [3] has also been introduced. Although all these terms can be viewed as synonyms, a particular term typically indicates a particular perspective, e.g., a physical distributed system perspective vs. a functional-oriented service perspective.

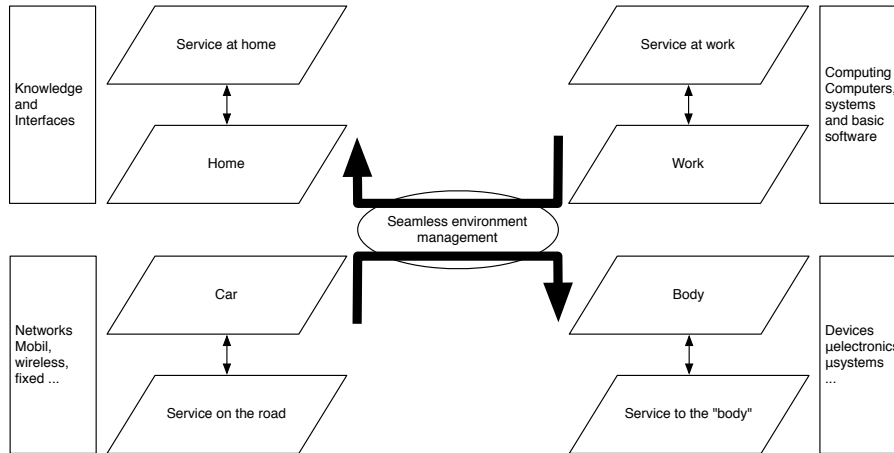


Fig. 1. The AmI space (adapted from [9])

the literature still seem like science fiction [8]. When inspecting the literature it appears that the problems of “grasping” everyday life is assumed solved by some sort of intelligent behaviour, but the methods that realise those behaviours are seldom addressed explicitly. The long term vision of the ISTAG group is the concept of *AmI space*, which comprises networked embedded systems that host dynamically configurable services (see Figure 1). The AmI space realises the vision of ambient intelligence by integrating local functionality across a variety of environments. This enables direct, natural and intuitive interaction between the user and services spanning a collection of environments.

In earlier work we have reported on the feasibility of case-based reasoning as a method for realising context-awareness in an ambient intelligent setting [10,11]. In particular, knowledge intensive case-based reasoning, in which the case-based reasoning process is supported by a structural model of general domain knowledge, appears promising. Case-based reasoning lends itself easily to reasoning about context and situations. The fact that case-based reasoning springs from work on understanding reasoning as an explanation process [12], and approaches reasoning by storing and remembering specific episodes (situations), clearly suggests it as a candidate for computational reasoning about situations.

The focus of the work presented here is on the situation awareness task, and more specifically on classifying the situation types in question, based on context elements such as location, time, or type of people involved. Our application domain is the hospital ward domain. The setting is that several times during the day medical personnel gather together for some purpose, with or without the presence of patients. The system’s task is to determine what type of gathering this is. An experimental system has been developed, within the framework of the CREEK system [13,14], and adapted for use on mobile pocket-size computers

suitable for being carried around by ward personnel. The system’s architecture, knowledge structures, and methods are presented in an earlier paper [10], which also exemplifies how a ward situation is classified. In the present paper we briefly review the system architecture and main components, and describe an evaluation study that has been made to assess the system’s quality. The system, called AmICREEK², is evaluated according to its ability to classify situations correctly [15]. The results are analysed and discussed within a larger evaluation framework that also include some of the evaluation criteria put forth by Cohen [16].

The rest of the paper is structured as follows: First a short overview of related work is presented. This is followed, in Section 3, by a summary of the architecture of AmICreek. In Section 4 the hospital ward study is presented, with the outcome analysed and discussed in Section 5. A conclusion and outlook on future work ends the paper.

2 Related Work

Traditionally, case-based reasoning has been applied to monolithic decision support systems. Recently, however, case-based reasoning has also been applied within the ambient intelligent community. Zimmermann [17] reported on case-based reasoning used to generate recommendations based on the user’s context in a mobile environment. The user context was encapsulated inside cases to facilitate comparison of contexts, generating recommendations based on case similarities, and learning of user behaviour. This work was part of a project concerned with audio augmentation of the real world in the context of the art museum in Bonn. Adapting solutions to particular users was also a focus in the work by Ma et al. [18], where case-based reasoning was used to adapt the behaviour of smart homes to users’ preferences. Multi-user smart homes can, even with a very limited amount of connected devices, present themselves with a very large amount of possible key processes and dependencies between them. Case-based reasoning was in this work used to identify these interdependencies, due to its ability to reason in ill understood and poorly structured domains. B  nard et al. [19] investigated the use of case-based reasoning as a mechanism for selecting suitable behaviour in different situations. They proposed an agent-based architecture that uses perceived information, or context, as the findings of a case and the proposed action as the solution. The authors ground their context model in a psychological framework that resembles the idea of selective interest and background context as argued by Dewey [20]. This is similar to the approach taken in earlier work within our group [21], which focused on the development of a comprehensive context model. As in that work, B  nard et al. separated context into external and internal context. The former includes the entities present in the environment, and the latter includes skills, states of the

² AmICREEK is a recent label for the system, which in earlier reports has been referred to merely as a CREEK system (referring to its framework and CBR method types), or a TrollCREEK system (referring to the Java implementation of CREEK which it was developed from).

agent, the agent's strategies and the agent's history. The existing cases in the case base are pre-classified situations modelled by a domain expert.

Kwon et al. [22] applied case-based reasoning in a multi-agent environment, to estimate the best purchase in comparative shopping. The goal of the system was to achieve a best possible solution between seller and buyer. This system was tested in simulation with several experiments. The three tests were: no negotiation, only price negotiation, and price as well as quality negotiation. As expected by the authors, the complex negotiation outperformed the other two, with respect to buyer's payoff, seller's payoff, and seller's rate of bid winning.

The MyCampus [23] system used case-based reasoning to learn a user's context-sensitive message filtering preferences. When a new message arrived it was compared to the existing messages stored in the case base. The new case was compared to the type of message, the sender of the message, the user's current calendar activity, the user's current location, and the weather. An experiment was carried out to validate the feasibility of the case-based approach. The experiment showed that the accuracy of the filtering process grew from 50 % to over 80 % when using the case-based reasoning approach.

The latter two methods were evaluated by assessing the quality of the best solution, determined by the accuracy of the solution value when compared to the expected value. This is a well-established approach [15], and also the assessment strategy adopted to assess the AmICreek system. As pointed out in [24], however, evaluating the quality of case-based reasoning systems may not be as straightforward, for example when cases are complex structures which are time consuming to construct, or when the user benefits from a larger part of a rich case than just a particular value. These concerns should also make some bells ring when building and evaluating user-interactive systems in health care. Extended approaches to evaluation of case-based reasoning systems include the work of Santamaria et. al [25] addressing the complexity of case-based reasoning applications, the work by McSherry [26] on diversity as a criterion, and the work by Smyth and McKenna [27] on case base competence. Cohen has addressed the evaluation of AI systems in general [28], as well as case-based reasoning systems in particular [16]. In addition to the common accuracy assessment, we relate our result to some of Cohen's criteria in the discussion chapter.

3 System Architecture

AmICREEK is an ambient intelligent system that spans a part of the AmI space described in Section 1. AmICREEK observes the ongoing work, thereby becoming *aware of the presence of a person*; assess ongoing situations, thus *perceives the needs* of this person, and *responds intelligently* to ongoing situations. The system is implemented as a three layer architecture where each layer has its own specific responsibility [10]. The AmICREEK architecture is comparable to the model of situation awareness as described by Endsley et al. [29].

The functional architecture is depicted in Figure 2. The main layer of interest in the work presented here is the Awareness layer.

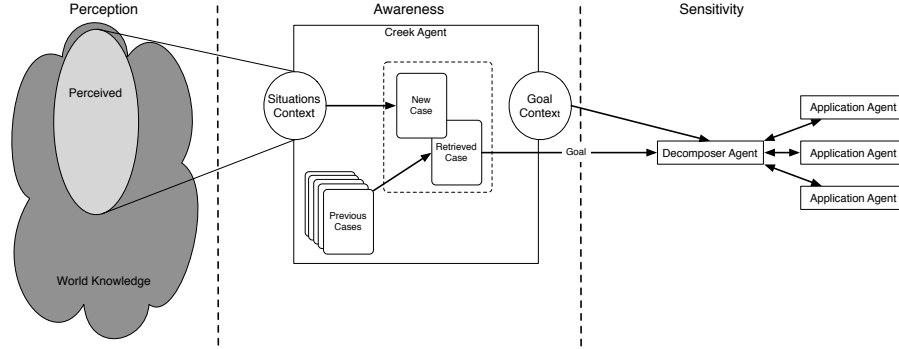


Fig. 2. Functional System Architecture

The initial layer is the *Perception* layer, which corresponds to the *Perception* layer in Endsley's model of situation awareness [29]. This layer is responsible for perceiving the environment and acquiring the necessary knowledge to feed the next two layers. The layer is implemented as a middleware solution that structures gathered information into a coherent structure [30], here named the *Situation Context* (see the middle part of Figure 2). Changes in the Situation Context will trigger an event that will initiate the case-based reasoning cycle.

The second layer is the *Awareness* layer. This corresponds to the second layer in Endsley's model (*comprehension*), and is the layer that the work presented here focuses on. This layer, which is the layer that exhibits the context-awareness, is implemented using the CREEK method [13,14]. The awareness layer acquires all relevant knowledge from the Perception layer and represents it as the findings of a new case [10]. The new case is matched to the existing case base, the goal is extracted from the best matching case and handed over to the third layer.

The third layer is the *Sensitivity* layer, which is comparable to the third layer in Endsley's model (*Projection*). This layer acquires the goal identified by the case-based reasoning cycle and constructs a sequence of tasks to execute that will satisfy the goal [31].

Context is used in two different ways within the architecture. Initially context describes the information that is perceivable in the world, labelled as Situation Context in Figure 2. This information constitutes the findings part of the cases, and is important for retrieving similar cases. The second use, labelled Goal Context in Figure 2, is the subset of the Situation Context that is relevant to the goal extracted from a matching case. For a more thorough description of the dualistic use of context and its place in the AmICREEK's knowledge model, please see [10].

Parameter	Description
Location	The room where the situation occurred
User	The user of the system
Role	The role of the user
Present	Other persons present
Role	The role of each of the persons present
Time	The time of day
Source	Information sources and targets
I/O	The direction of the information flow
Information	Type of information
Situation	The type of situation

Table 1. Essential Aspects of Situations

4 A Hospital Ward Study

The test presented here tests the Awareness layer, by focusing on the ability of the case-based reasoning system to assess situations correctly. The perception layer and sensitivity layer have been tested separately on other occasions. The reader is directed to [32] for further details.

4.1 Test Setup

The essential aspects of situations, in the hospital ward domain, are the nine parameters shown in Table 1. These parameters were noted in the ethnographical study detailed in [11]. The topmost six parameters are used to describe the situation as a case and constitute the findings. The next three are used to describe how a goal can be decomposed. Finally, the last parameter is the situation type as classified by the hospital personnel.

For the initial test, the data observed by a human observer, following consultant physician number nine (OL9) were chosen. The physician was the one with the highest number of experienced situations, as well as the one who experienced the highest number of different situations. For the purpose of this test, two days (13 and 14) of observations of OL9 were used. The evaluation of the system's performance was initially done by a qualitative evaluation of the data from the cardiology ward. This was carried out in order to review the context model and the integration of the knowledge model. It is worth noticing that attending one meeting can be viewed as many situations (in sequence). For example, a given meeting of the type pre-ward-round will typically involve the discussion of several patients. Thus, such a meeting is broken down into one pre-ward-round situation per patient discussed.

4.2 Test Execution

The test was conducted as a three-step process. First the 25 different situations that occurred on day 13 were added to the case base, partly as solved and un-

Situation type	Number of situations	Percentage of all situations covered by day 13	Initial number of solve cases	Sample
Post-work	9	69%	6	1,2,3,4,6,9
Preparation	1	25%	1	1
Pre-ward-round	6	23%	4	2,3,4,6
Patient meeting	1	25%	1	1
Examination	2	25%	1	1
Ward-round	6	23%	4	3,4,5,6

Table 2. Distribution of Observed Situations for OL9, day 13

solved cases, and the retrieve-step was executed for each of the unsolved cases. Secondly, the learning process of the case-based reasoning algorithm was manually executed. That is, the former unsolved cases were classified and flagged as solved. Finally, all of the day 13 cases were used to classify the situations occurring on day 14.

4.3 Initial Modelling

To get a representative distribution of solved and unsolved cases, the 25 different situations for day 13 were divided into two groups. $\frac{2}{3}$ of the cases (17) were randomly selected and marked as solved, the remaining $\frac{1}{3}$ (8) were added as unsolved cases. Table 2 describes the number of different situations, the percentages of the total situations, how many were initially selected as solved, and the specific situations in question.

The column labelled *Situation type* describes the six different situations experienced by physician OL9. *Number of situations* counts the number of each situation-type occurring on day 13. The column labelled *Percentage of all situations covered by day 13* shows us how large a percentage the specific situation type that occur on day 13. As an example, the nine observed situations of the type *Post-work* covers 69 % of all observed situations of this type for both day 13 and 14. The column labelled *Initial number of solved cases* gives us the number of cases chosen to be marked as solved. Finally, the column labelled *Sample* tells us which specific situations were modelled as solved cases.

The situations chosen as solved were modelled as cases using the six upper parameters shown in Table 1 as findings. The parameter *Situation* was used to mark each of the cases with a top-level goal, in the sense that for e.g. the *Post-work* situations the goal *Post-work-goal* was identified.

After the modelling of the cases marked as solved the remaining cases were modelled as unsolved. Each of these case were modelled in exactly the same fashion as those solved, except for obviously being marked as unsolved. Each of the unsolved cases were then run through the retrieve-step and matched to the solved cases.

Input case	Strength Matching case	
Ward round 1301	89%	Pre-ward round 1306
	88%	Pre-ward round 1304
	88%	Pre-ward round 1303
Ward round 1302	88%	Ward round 1305
	88%	Pre-ward round 1306
	88%	Pre-ward round 1304
Pre-ward round 1301	100%	Pre-ward round 1303
	100%	Pre-ward round 1304
	99%	Pre-ward round 1306
Pre-ward round 1305	100%	Pre-ward round 1304
	100%	Pre-ward round 1306
	100%	Pre-ward round 1303
Examination 1302	100%	Examination 1301
	55%	Post work 1304
	54%	Post work 1303
Post work 1305	100%	Post work 1306
	65%	Preparation 1301
	64%	Post work 1304
Post work 1307	99%	Post work 1306
	66%	Preparation 1301
	64%	Post work 1309
Post work 1308	66%	Post work 1309
	65%	Preparation 1301
	61%	Consultation 1304

Table 3. Result of Matching Test (Run 1)

4.4 Executing Day 13 Test

Table 3 demonstrates the result of the retrieve process for the unsolved cases of day 13. The current implementation only selects the best matching case for goal extraction. So when examining the cases in Table 3 that are classified correctly we can disregard all but the best matching case. Following that line of reasoning only **Ward round 1301** has been misclassified. For **Ward round 1302**, **Ward round 1305** gets selected due to the detailed mechanism for ranking matched cases not apparent in Table 3. AmICREEK suggests that this particular ward round is a pre-ward round, and that the best matching case is **Pre-ward round 1306**.

If we examine the two cases shown, see Figure 3, we will see that they do not look like they should resemble each other to such a high degree as suggested. The differences between the two cases have been marked with dashed-lined circles. Starting with the **Environmental Context**, it contains **Patient #36**, which the matching case does not. Obviously this missing feature should suggest that the two cases differ. The CREEK method does support the ability to lower the matching strength when a matching case misses features. However, the AmICREEK does not yet implement this feature. Looking at the two time

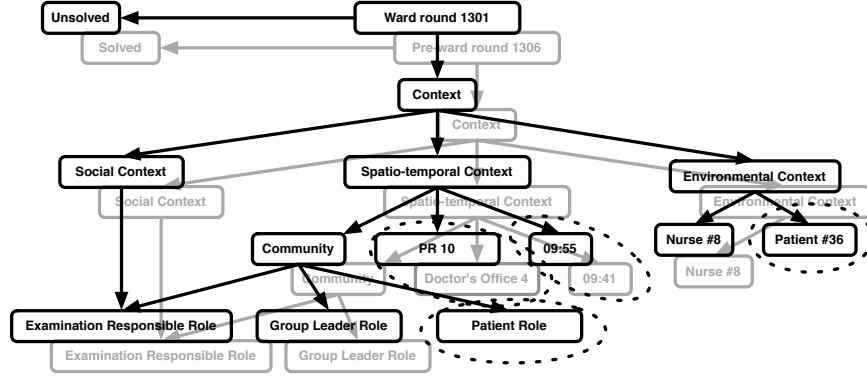


Fig. 3. Ward round 1301 vs. Pre-ward round 1306

values, we can observe that they are very close, and are indeed reported as being 98 % similar. The unknown case occurs in the patient room PR10, whereas the best matching case takes place in the **Doctor's office 4**. AmICREEK decides, correctly, that these two locations do not syntactically match. However, both of these locations are instances of the general class **Location**, and as the knowledge model is structured in a semantic network it should be possible to explain that they are somewhat alike. The CREEK method does allow for this type of reasoning by calculating the convergence point between two concepts [14]. However, currently AmICREEK has been focused on calculation of convergence points between causal relations, thus the calculation for non-causal relations, e.g., *instance of* and *subclass of*, results in convergence between all concepts, as all concepts at some point are an instance of, or subclass of the top-most concept *Thing*. Finally, the unknown case contains a **Patient Role**, not found in the best matching case. This is the same problem as described above regarding the patient in the **Environmental Context**.

If we examine the case base for known cases of the type **Ward round**, the best matching case is **Ward round 1303**, which AmICREEK reports as matching with a strength of 83 %, and as number seven in the list of matching cases. If we examine that case we will discover that it resembles the unknown case a lot. Only two parameters separates them: the time is 09:55 versus 10:13, something which AmICREEK calculates as a 91 % match; and the fact that the patient present is not **Patient #36** but **Patient #38**. It would seem reasonable that these two cases had a high matching strength, however as described above, other cases do not get their matching strength lowered, thus better matching cases cannot compete. This becomes quite evident if we change the patient present to **Patient #36**, in that case the matching strength between the two cases are 99 %, easily surpassing the matching strengths of the cases of the wrong type.

Input case	Strength Matching case	
Pre-ward round 1402	100%	Pre-ward round 1303
	100%	Pre-ward round 1301
	100%	Pre-ward round 1304
Pre-ward round 1404	100%	Pre-ward round 1305
	100%	Pre-ward round 1306
	100%	Pre-ward round 1304
Pre-ward round 1407	100%	Pre-ward round 1306
	100%	Pre-ward round 1305
	100%	Pre-ward round 1304
Examination 1404	55%	Examination 1302
	55%	Examination 1301
	54%	Post work 1304
Examination 1405	54%	Examination 1302
	54%	Examination 1301
	54%	Preparation 1301
Ward round 1402	88%	Ward round 1305
	88%	Pre-ward round 1306
	88%	Pre-ward round 1305
Ward round 1405	99%	Ward round 1305
	99%	Ward round 1304
	89%	Pre-ward round 1302
Ward round 1408	89%	Pre-ward round 1302
	88%	Pre-ward round 1301
	87%	Ward round 1305

Table 4. Result of Matching Test (Run 2)

4.5 Executing Day 14 Test

As described above, Table 3 shows how AmICREEK classified the cases from day 13 that had be flagged as unsolved. To extend this simulation, we can now move into day 14 by manually executing the learning process, in other words correctly classify these previously unsolved cases and flag them as solved. Table 4 depicts the results of attempting to match the situations from day 14 that was marked as unknown against all the situations from day 13.

If we examine Table 4 we will see that most of the cases were classified correctly. Only **Ward round 1408** is classified wrongly, as a **Pre-ward round**. Looking at **Ward round 1408** it matches two pre-ward rounds before **Ward round 1305** is matched with a matching strength of 87 %. Again the explanation lies in the fact that semantic matching is currently not implemented and cases are not being “punished” for missing parameters. In this case, the fact that **Pre-ward round 1302** is occurring at 11:10 and **Ward round 1408** at 11:30 is the one parameter that forces the ranking of the pre-ward round as the best match. As with the example depicted in Figure 3, the pre-ward round neither contains a patient nor a patient role, and the location is wrong.

Situation type	Run 1		Run 2	
	best	all	best	all
Post-work	100%	55%	N/A	N/A
Pre-ward-round	100%	100%	100%	100%
Examination	100%	33%	100%	66%
Ward-round	50%	16%	100%	50%

Table 5. Absolute Accuracy of Case-Based Reasoning Test

4.6 Accuracy of the Classifications

The data gathered at the cardiology ward included an *in situ* classification carried out by the physician observed [11]. When measuring the accuracy of the classifications done in AmICREEK we start off from the assumption that these classifications are correct. Thus, we can compare AmICREEK’s classifications to those carried out by an expert. If we examine Table 5 we can see the absolute accuracy for the four different types of unknown situations that AmICREEK attempted to classify. For both runs the table gives the percentage of times that the best matching case was correct, and the percentage of cases that was correct within the top-three across all matches of this type. The latter can be regarded as the confidence of the classification. Thus, a low number will tell us that the classification was distributed, whereas a high number will tell us that the classification was uniform. As an example, for the **Post-work** cases in *Run 1* all the best matching cases were correct. However, only 55 % of all top-three matches were correct.

As shown in Table 5, AmICREEK does a reasonably good job of classifying the situations correctly. **Post-work**, **Pre-ward-round** and **Examination** are all classified correctly in *Run 1*, and all were classified correctly in *Run 2*³. Yet, **Ward-wound** situations were only classified correctly 50 % of the time in *Run 1*.

If we look at the distribution among the top-three matches for each situation type we can see that except for situations of the type **pre-ward-round**, most did not classify perfectly. However, as aforementioned, goal extraction was only carried out for the best matching case.

5 Analysis and Discussion

The evaluation presented is an evaluation of case-based reasoning as a method for achieving situation-awareness. The use of case-based reasoning as the method of classification is rooted, not only in the desire to find the best suited algorithm for a limited data set, but also in the fact that the theory on human cognition that underlies case-based reasoning is one of situation classification. Thus, the cognitive plausibility [33] of using case-based reasoning is to be found in the

³ There were no situations of the type **Post-work** occurring on day 14

general characterisation of case-based reasoning as a way of human reasoning. Hence, the metric for evaluating the use of case-based reasoning in our work, is not its cognitive validity [16], but rather the performances of the classification.

Classification performance can be evaluated in four ways [15]: *i*) absolute accuracy, as determined by a domain expert; *ii*) the plausibility of incorrect classification, also determined by an expert; *iii*) performance can be compared to other algorithms; and *iv*) performance can be compared to that of an expert.

Currently we have no other implementation of an alternative algorithm that fits the data readily available, thus no comparison is available (part *iii* above). Further, as stated above the choice of case-based reasoning as the algorithm of choice is not purely a performance choice, but also one of origin. In addition, the performance of AmICREEK is implicitly already comparable to that of a human expert (part *iv* above). The subjects observed did already do classification during data collection [11], and the classification was *in situ*, thus we must assume that these classifications are correct. For the remaining part of the discussion we focus on the accuracy of the classification and plausibility of incorrect classifications.

Regarding the plausibility of misclassifications, it could be argued that to determine if any misclassifications are plausible an expert working at the particular hospital ward is required. However, given the nature of the observed data, and the amount and distribution of parameters, it seems reasonable that the same assumption regarding the expert with respect to accuracy also holds for misclassifications.

Revisiting the results of the two runs described in Section 4 in light of absolute accuracy, we can re-examine Table 5, we can summarise the accuracy as very good. All possible classifications, except the situation type **Ward-round** were classified correctly. Thus, the absolute accuracy is all in all very good compared to the expert who classified the situations when they were observed.

If we investigate further, we can see that some of the classifications in the top-three matches were not correct. In particular with respect to the situation type **Ward-round**, the situation was misclassified half of the time in *Run 1*. To defend the performance of AmICREEK an explanation for this incorrect classification must be given. Section 4 did already include explanations as to why each of the misclassified cases was classified as it was. The two main general reasons are as follows:

Unsolved cases do sometimes contain features not found in the retrieved case. These extra features should lower the matching strength of the retrieved case. However, the particular implementation of how the matching strength is calculated does currently not take this into account. Thus, the retrieve process uses incorrect matching strengths with respect to this effect. These extra features are very obvious to a human expert, and disregarding them goes a long way in explaining why some cases were misclassified.

Secondly, when a feature is matched in the leaf nodes of the case structure, symbols such as **Location**, **Person** or **Role** are only matched on the syntactical level. This means that the fact that the instances of these concepts are located in the multi-relational semantic network is disregarded in the retrieve process.

As with the problem above, the fact that, for an example, a patient is present is a strong indication for a domain expert that a **Ward-round** is taking place. The fact that support for semantic matches of non-causal relations are not yet implemented does explain why some cases were misclassified.

Finally, on day 13 only two **Ward round** cases were initially classified as solved, yet as more cases were solved in day 14 the classification dramatically improved (compare Table 3 and Table 4).

In summary, the situations are being classified at an acceptable level for this first experiment compared to the human domain expert. As we must assume that the classification gathered are correct, the human domain expert does slightly outperform AmICREEK. However, some, if not all of the misclassification can be explained by the way matching strengths are currently being calculated

6 Conclusion and Further Work

The work presented here demonstrates that knowledge intensive case-based reasoning is a promising method for constructing situation-awareness systems. By using the CREEK method, case-based reasoning is applied as the reasoning mechanism to identify ongoing situations. The assessment of situations is the very core of the intelligent behaviour of an ambient intelligent system. It has been demonstrated that using a knowledge intensive case-based reasoning methodology, where the cases are submerged into the general knowledge model, facilitates situation-awareness.

However, it has also been shown that some deficiencies in the current implementation of the CREEK method still exists. In particular with regard to *plausible inheritance*, which currently only works with causal concepts, and with the similarity function that currently does not calculate similarity correctly when cases are missing parameters. Remedying these two issues should significantly improve the system.

Finally, the results should be reproduced using publicly available open source systems. Primarily to verify that it is not the idiosyncrasies of our implementation that allows the reasoning to work, and secondly to allow others to freely use and scrutinise the system.

Acknowledgements

We would like to extend our gratitude to many of colleagues for their input throughout this work, and in particular to Jörg Cassens, without whom this work would not have been possible.

References

1. Weiser, M.: The computer for the 21st century. Scientific American (September 1991) 94–104

2. Hansmann, U., Merk, L., Nicklous, M.S., Stober, T.: *Pervasive Computing: The Mobile World*. Springer Professional Computing (2003)
3. Greenfield, A.: *Everyware: The Dawning Age of Ubiquitous Computing (Voices That Matter)*. New Riders Publishing (2006)
4. Lugmayr, A.: The future is 'ambient'. In Creutzburg, R., Takala, J.H., Chen, C.W., eds.: *Proceedings of SPIE*. Volume 6074 of *Multimedia on Mobile Devices II*, SPIE (2006)
5. Ducatel, K., Bogdanowicz, M., Scapolo, F., Leijten, J., Burgelman, J.C.: *ISTAG scenarios for ambient intelligence in 2010*. Technical report, IST Advisory Group (2001)
6. Aarts, E.H.L., Encarnação, J.L., eds.: *True Visions: The Emergence of Ambient Intelligence*. Springer (2006)
7. Satyanarayanan, M.: A catalyst for mobile and ubiquitous computing. *IEEE Pervasive Computing* **1**(1) (January 2002) 2–5
8. Lueg, C.: Representation in pervasive computing. In: *Proceedings of the Inaugural Asia Pacific Forum on Pervasive Computing*. (2002)
9. ISTAG: IST advisory group, strategic orientations and priorities for IST in FP6 (June 2002)
10. Kofod-Petersen, A., Aamodt, A.: Contextualised ambient intelligence through case-based reasoning. In Roth-Berghofer, T.R., Göker, M.H., Güvenir, H.A., eds.: *Proceedings of the Eighth European Conference on Case-Based Reasoning (ECCBR 2006)*. Volume 4106 of *Lecture Notes in Computer Science*, Ölüdeniz, Turkey, Springer Verlag (September 2006) 211–225
11. Cassens, J., Kofod-Petersen, A.: Using activity theory to model context awareness: a qualitative case study. In: *Proceedings of the 19th International Florida Artificial Intelligence Research Society Conference, Florida, USA, AAAI Press (May 2006)* 619–624
12. Schank, R.: *Dynamic memory; a theory of reminding and learning in computers and people*. Cambridge University Press (1982)
13. Aamodt, A.: A knowledge-intensive, integrated approach to problem solving and sustained learning. PhD thesis, University of Trondheim, Norwegian Institute of Technology, Department of Computer Science (May 1991) University Microfilms PUB 92-08460.
14. Aamodt, A.: Knowledge-intensive case-based reasoning in CREEK. In Funk, P., Calero, P.A.G., eds.: *Advances in case-based reasoning, 7th European Conference, ECCBR 2004, Proceedings*. (2004) 1–15
15. Bareiss, R.: The experimental evaluation of a case-based learning apprentice. In: *Proceedings of a Workshop on Case-Based Reasoning, Pensacola Beach, Florida, USA, Morgan Kaufmann (1989)* 162–167
16. Cohen, P.R.: Evaluation and case-based reasoning. In: *Proceedings of a Workshop on Case-Based Reasoning, Pensacola Beach, Florida, USA, Morgan Kaufmann (1989)* 168–172
17. Zimmermann, A.: Context-awareness in user modelling: Requirements analysis for a case-based reasoning application. In Ashley, K.D., Bridge, D.G., eds.: *ICCBR 2003, Case-Based Reasoning Research and Development*. Number 2689 in *Lecture Notes in Artificial Intelligence*, Springer-Verlag (2003) 718–732
18. Ma, T., Kim, Y.D., Ma, Q., Tang, M., Zhou, W.: Context-aware implementation based on CBR for smart home. In: *Wireless And Mobile Computing, Networking And Communications, 2005. (WiMob'2005)*, IEEE, IEEE Computer Society (August 2005) 112–115

19. Bénard, R., Bossard, C., Loor, P.D.: Context's modelling for participative simulation. In: Proceedings of the 19th FLAIRS Conference, AAAI Press (May 2006) 613–618
20. Dewey, J.: Context and Thought. University of California Press (1931)
21. Öztürk, P., Aamodt, A.: A context model for knowledge-intensive case-based reasoning. *International Journal of Human Computer Studies* **48** (1998) 331–355
22. Kwon, O.B., Sadeh, N.: Applying case-based reasoning and multi-agent intelligent system to context-aware comparative shopping. *Decision Support Systems* **37**(2) (May 2004) 199–213
23. Sadeh, N., Gandon, F., Kwon, O.B.: Ambient intelligence: The mycampus experience. Technical Report CMU-ISRI-05-123, Carnegie Mellon University (July 2005)
24. Gu, M., Aamodt, A.: Evaluating cbr systems using different data sources: a case study. In Roth-Berghofer, T.R., Göker, M.H., Güvenir, H.A., eds.: Proceedings of the Eighth European Conference on Case-Based Reasoning (ECCBR 2006). Volume 4106 of Lecture Notes in Computer Science., Ölüdeniz, Turkey, Springer Verlag (September 2006)
25. Santamaria, J.C., Ram, A.: Systematic evaluation of design decisions in cbr systems. In: Proceedings of the AAAI Case-Based Reasoning Workshop, Seattle, Washington, USA (1994) 23–29
26. McSherry, D.: Diversity-conscious retrieval. In Craw, S., Preece, A., eds.: Proceedings of the 6th European Conference on Case-Based Reasoning (ECCBR02). Volume 2416 of Lecture Notes in Artificial Intelligence., Springer-Verlag (2002) 219–233
27. Smyth, B., Mckenna, E.: Modelling the competence of case-bases. In: Advances in Case-Based Reasoning: Proceedings of the Fourth European Workshop on Case-Based Reasoning, Springer-Verlag (1998) 208–220
28. Cohen, P., Howe, A.: How evaluation guides ai research. *AI Magazine* (9) (1988) 35–43
29. Endsley, M.R., Bolté, B., Jones, D.G.: Designing for Situation Awareness: An Approach to User-Centered Design. Taylor & Francis (2003)
30. Kofod-Petersen, A., Mikalsen, M.: Context: Representation and reasoning – representing and reasoning about context in a mobile environment. *Revue d'Intelligence Artificielle* **19**(3) (2005) 479–498
31. Gundersen, O.E., Kofod-Petersen, A.: Multiagent based problem-solving in a mobile environment. In Coward, E., ed.: Norsk Informatikkonferance 2005, NIK 2005, Institutt for Informatikk, Universitetet i Bergen (November 2005) 7–18
32. Kofod-Petersen, A.: A Case-Based Approach to Realising Ambient Intelligence among Agents. PhD thesis, Department of Computer and Information Sciences, Norwegian University of Science and Technology (2007)
33. Cohen, P.R., Howe, A.E.: How evaluation guides AI research. *AI Magazine* **9**(4) (1988) 35–43