Explanations and Case-Based Reasoning in Ambient Intelligent Systems

Jörg Cassens and Anders Kofod-Petersen

Department of Computer and Information Science (IDI),
Norwegian University of Science and Technology (NTNU),
7491 Trondheim, Norway,
{cassens|anderpe}@idi.ntnu.no,
http://www.idi.ntnu.no/

Abstract. Interacting with intelligent systems in general and ambient intelligent systems in particular, requires that these systems have the ability to build a trust relationship with the users. The ability to explain its own behaviour is one of the most important abilities that such a system can exhibit to gain trust. We argue that explanations are not just an addition to an ambient intelligent system rather it is an approach to the design and implementation of such a system. Explanations are useful both for the reasoning process itself and as a means of communicating with the users. In this paper, we present a knowledge intensive approach for identifying different contexts and generating a course of action depending on the context found. We explore the use of explanations both as a means of reasoning and as a means of communication with the user.

1 Introduction

Ambient intelligent systems are described by their ability to be aware of its users, perceive their needs and respond intelligently to them [1]. To achieve this, such a system must exhibit pro-activity and reasoning. Thus, contrary to traditional systems where the user is in charge and the system plays a passive role, ambient intelligent system can assume responsibility and behave proactively. The shift from passive systems, to what could be regarded as a partnership between humans and intelligent artefacts, fosters the need for social adept system [2], in such a way that intelligent systems have to show certain abilities traditionally ascribed to humans [3]. Among these abilities, we would count a system's ability to explain its behaviour.

Further, the nature of an ambient intelligent system dictates that the way we traditionally interact with computer systems are substituted with multi modal interfaces and in particular behavioural interfaces. It is not only the input devices where behavioural interfaces are used; the main output device for an ambient intelligent system is its behaviour. Thus, a user now runs the risk of having to master the methods of behavioural psychology to use an ambient intelligent system. Relying solely on observing the behaviour of an ambient intelligent system will, at best, give us a limited understanding of its behaviour. It will not allow us to consider any internal organisation. Yet, to persuade the user of a system's usability and credibility, explanations are in order

The work presented here argues that explanations are not just an addition to an ambient intelligent system; rather it is an approach to the design and implementation of such a system. Explanations are a useful approach to both the reasoning process from a system centric perspective, as well as a means of communication viewed from a user centric perspective.

This paper is structured as follows: First, an overview on how explanations and case-based reasoning are related is given. Secondly, an introduction to the use of case-based reasoning in ambient intelligence, as well as a short overview of the system employed here, is presented. Thirdly, a description of the theoretical framework underpinning our approach to the use of explanations and case-based reasoning in ambient intelligence is given. A summary and pointers to future work ends the paper.

2 Explanations and Case-based Reasoning

Originally, case-based reasoning did emerge from an understanding of reasoning as an explanation process [4,5]. Schank describes explanations as the most common method used by humans to support their decision making. An understanding of common occurrences assists in comprehending stories; in such a way that details omitted or assumed implicitly do not make a story incomprehensible.

Sørmo et al. [6] present a framework for explanations in intelligent systems with a special focus on case-based reasoning. Specifically, they identify five goals that explanations can satisfy. The goal of *transparency* is concerned with the system's ability to explain how an answer was reached. *Justification* deals with the ability to explain why the answer is good. When dealing with the importance of a question asked, *relevance* is the goal that must be satisfied. *Conceptualization* is the goal that handles the meaning of concepts. Finally, *learning* is in itself a goal, as it teaches us about the domain in question.

Roth-Berghofer has explored some fundamental issues with different useful kinds of explanation and their connection to the different knowledge containers of a case-based reasoning system [7]. Five different kinds of explanation are identified: *conceptual explanations*, which map unknown new concepts to known ones, *why-explanations* describing causes or justifications, *how-explanations* depicting causal chains for an event, *purpose-explanations* describing the purpose or use of something, and *cognitive explanations* predicting the behaviour of intelligent systems. Roth-Berghofer further on ties these different kinds of explanation to the different knowledge containers of case-based reasoning systems [8], namely case base, similarity measure, adaptation knowledge, and vocabulary.

Building on the last two works, we have earlier started to investigate a combined framework of user goals and explanation kinds [9]. The goal of this work was to outline a design methodology that starts from an analysis of usage scenarios in order to be able to identify possible expectations a user might have towards the explanatory capabilities of an intelligent system. The requirements recognised can be used to identify which kind of knowledge has to be represented in the system, and which knowledge containers are best suited for this task. In this work, we have identified the need for socio-psychological analyses of workplaces in order to be able to design systems which can meaningful engage in socio-technical interactions. The advantages of designing systems from a socio-technical perspective has been investigated through the use of activity theory as a method for designing an ambient intelligent case-based reasoning system [10].

3 Case-based Reasoning and Ambient Intelligence

Weiser, who coined the term *ubiquitous computing*, did explicitly state that artificial intelligence was unimportant when realising the visions described [11, p.3]. This view that proper embedding of computers is sufficient does also to a large degree saturate the field of *pervasive computing*. However, the insight that achievement of the visions described in pervasive computing does indeed require some degree of reasoning has lead to the developments jointly labelled as *ambient intelligence*.

The IST Advisory Group to the European Commission (ISTAG) describes ambient intelligence as human beings surrounded by intelligent interfaces, supported by computing and network technology embedded in everyday objects. More importantly, the environment should be aware of the presence of a person, perceive the needs of this person and respond intelligently to them [1].

The ability to be aware of the environment, reason about ongoing situations and decide about appropriate behaviour is closely linked with being knowledgeable about the world. Case-based reasoning in general, and knowledge intensive case-based reasoning in particular [12,13], appears to be a promising candidate for reasoning about situations and behaviour in an ambient intelligent setting.

Zimmermann reports on case-based reasoning used to generate recommendations based on the user's context in a mobile environment [14]. The user context is encapsulated inside cases to facilitate comparison of contexts, generating recommendations based on case similarities, and learning of user behaviour. The cases are structured around a context describing the user's environment as the findings and a recommendation for a particular audio file as the solution.

Along the same lines of adapting solutions to users is the work by Ma et al., where case-based reasoning is used to adapt the behaviour of smart homes to user preferences [15]. The cases are represented as frames, where the findings are: the user, the environment, the time, and the values of the active devices. When new cases are instantiated; the similarities between the new and existing cases are calculated and a setting for the appliances in the house is selected.

Recently Bénard et al. have investigated the use of case-based reasoning as a mechanism for selecting suitable behaviour in different situations [16]. They propose an agent-based architecture that uses perceived information, or context, as the findings of a case and the proposed action as the solution. The existing cases in the case base are pre-classified situations modelled by a domain expert. The authors briefly describe how the case-based reasoning process approaches

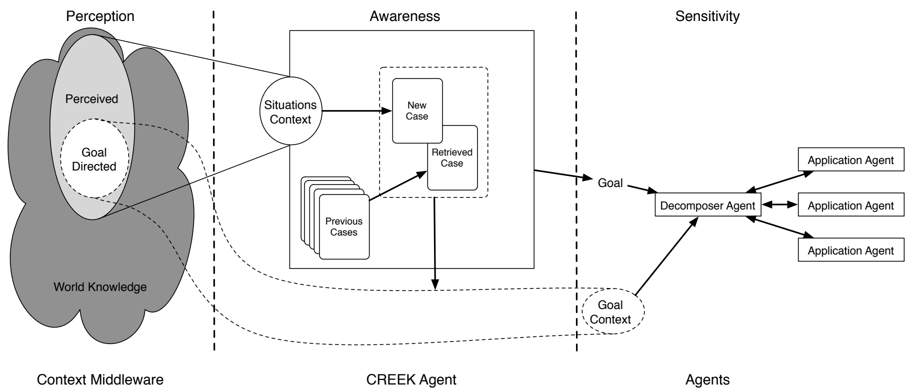

Fig. 1. Functional System Architecture

the problem of behaviour selection. Initially a new context is compared to the findings in the existing cases and the best matching case is retrieved. Secondly, the retrieved case is adapted to fit the new context.

The work presented here builds on the experience gained by applying case-based reasoning to situation assessment in an ambient intelligent system [17]. The architecture developed and system implemented approaches ambient intelligence by separating the main responsibilities into three layers.

The architecture has been implemented as an ambient intelligent system in a hospital ward [18]. The personnel at the hospital ward are involved in many different activities, such as doing ward rounds, meetings and different forms of examinations. The system's main purpose is to recognise ongoing situations and proactively acquired digital information relevant for the user.

Figure 1 depicts the functional system architecture. The Context Middleware [19] provides an infrastructure that perceives the environment by collecting and maintaining a coherent model of available context. The CREEK agent implements and extends the knowledge intensive case-based reasoning method CREEK [13], which is responsible for assessing occurring situations through context awareness. This is done by constructing an unsolved case where the findings are the context received from the Context Middleware. This case is then matched to existing cases and the best matching case contains the goal for this particular type of situation. This goal is then transmitted to the Decomposer Agent, which decomposes according to the existing artefacts and persons described in the context. Each of these entities are described by an Application Agent that supplies information. These two types of agents are together responsible for executing suitable system behaviour by being context sensitive.

The knowledge model utilised in this system is structured as a multi-relational semantic network, where each of the five different parts are integrated. Three parts of the knowledge model are universal for any CREEK application: the top-level ontology, called ISOPOD; the domain-specific model of general and factual knowledge; and the case base. In addition to these, two parts have been developed for the ambient intelligent system, namely the Basic Context Model and the Activity Theory Model.

The Basic Context Model is structured around a meronomy based on tradition in pervasive computing [19,20], which imposes a structure that facilitates easy development of context sensitive applications [21]. The Activity Theory Model captures knowledge regarding activities, which is one of the most important aspect of situations [22].

4 Facets of Explanation and Context

As outlined in Section 2, the term explanation can have different focii in artificial intelligence. In his introduction to the ECCBR-04 Workshop on Explanation in case-based reasoning, David Leake [23] has identified three different facets of explanation and case-based reasoning:

- Using explanations to support the case-based reasoning process
- Generating explanations by case-based reasoning
- Using cases for explaining system results to an external user

With our notion of user goals, we can subsume the last two facets as both being targeted towards the user of the system. In our understanding, showing the case to the user is a special case of 'generating explanations by case-based reasoning', making use of the case-based reasoning assumption that similar problems have similar solutions. Provided that the user has some knowledge about the similarity function and that the case structure is understandable by the user, the displayed case acts as an explanation to the user (see e.g. [6,24]). We are left with two functions of an explanation, as described in [25]: first, enhancing and promoting the reasoning process. Second, delivering some knowledge about the reasoning process, its results, or implication to the user. We call the first aspect the *system centric view* on explanation and the second one the *user centric view* on explanation:

- Explanation as **part of the reasoning process** itself.
  Example: a knowledge intensive case-based reasoning system can use its domain knowledge to explain the absence/variation of feature values.
- Giving explanations of the found solution, its application, or the reasoning process **to the user**.
  Example: in an engine failure diagnosis system, the user gets an explanation on why a particular case was matched.

When we now look at the use of context, we can again identify different views on context during the reasoning process and the use of the system. First, the system has to identify an unknown situation out of its internal state and the percieved state of the world. Here the system has to find out which situation it

Table 1. Explanations

	Context Awareness	Context Sensitivity
System Centric	Generate an explanation to **recognise** the situation	**Identify** the behaviour the system should expose
User Centric	**Elucidate** why the system identifies a particular situation	**Explicate** why a certain behaviour was chosen

is in. Second, the system has to act according to the percieved state of the world and an assumed context. Following our earlier argumentation introduced in [18], we identify these two aspects as two distinct steps in the reasoning process:

- **Context Awareness:** Trying to figure out which situation the system is in. *Example:* An ambient intelligent system for supporting health personnel recognises that the user is on a ward-round because of the time of the day, the location, and the other persons present.
- **Context Sensitivity:** Acting according to the situation the system thinks it is in.
  Example: the same system fetches the newest versions of electronic patient records of all patients in the room from the hospital systems. When the user stands close to the bed of a patient, the system automatically displays them.

Combing these views on explanation and on context, we end up with two dimensions of inquiry as depicted in Table 1.

5 Using Explanations in an Ambient Intelligent System

We will further investigate the relationship between context and explanations by examining an example from our current implementation. At this time, our system is capable of assessing ongoing situations in a simulated hospital (cardiology) ward domain. The system's main purpose is to identify ongoing situations and proactively acquire digital information required by the persons present.

Recognise: In this step, we are *using explanations to recognise the current ongoing situation*. The system uses all available resources in its reasoning process. Let us assume that the ongoing situation is a *ward round*. Normally ward rounds take place in a patient room, however the current situation is occurring in the hallway. This discrepancy can be explained away by the system generalising that both locations can indeed contain patient beds. When CREEK retrieves a matching case, the system has no explicit knowledge stating that a hallway can contain hospital beds. The initial match is of a syntactical nature only. However, it can use its general knowledge and the reasoning mechanism of plausible inheritance to generate an explanation supporting the hypothesis that beds can be located in the hallway, for example because they are both some kinds of rooms, and beds are some kinds of objects which have a room as a location. Therefore, as all other parameters are consistent with a ward round, the system assumes that it is indeed a ward round situation.

The explanation used by the system in this example states that a hallway is a room and can therefore contain a hospital bed.

Elucidate: We now want to *generate an explanation for the user that tells the user why the system assumes a certain situation*. The system will make use of all available sources of knowledge in order to gain the user's confidence in its capabilities. It will also have to consider the user's goals when chosing a specific explanation. It has been shown that simply presenting the reasoning trace is not always sufficient (it can even be counter productive) [26,27]. The system might therefore generate an after-the-fact explanation, which for example justifies its assumption instead. Since the ward round situation is occurring at an unusual place, CREEK will point out the time of the day, the availability of the other expected participants, and the fact that hallways might contain beds, as the reason for its assumption instead of only displaying its generalisation of the location.

The explanation shown to the user is a justification of the system's believe of being on a ward round.

Identify: After the system has successfully identified the context, it is *using explanations to generate a plan for a reasonable course of action*. Now, it is using only the knowledge sources important for the situation at hand (the context is acting as a focus lense [18]). When we now presume that the system has recognised that we are on a ward round, discussing medical conditions and treatments with several patients, it will try to prepare all the relevant information to be presented to the user. This includes all test results. The system can now ask other available artefacts for test results on the user, and the medical images database can offer a MR image whereas the patient record offers a textual description of the MRI. Because of limitations of handheld devices, the system will for example not be able to display high resolution MR images. When choosing which of the artefacts to query, the system will reject the medical image database and only query the electronic patient record database.

The explanation used by the system is based on the knowledge that a high resolution image displaying device is not available on a ward round.

Explicate: Looking at the user centric part again, we are now in need of *generating an explanation for why the systems takes a specific action*. The system will take into account which situation it assumes it is in and the possible goals the user might have for an explanation. In executing its plan, the system proposes its user to visit the isolation room with patients which should be kept seperate. The user is surprised since he is not aware that any of the patient he should see at the ward round is in the isolation room, and no information on this was exchanged in the morning briefing. The system can then generate an explanation that shows the relevance of the proposal by pointing out that one particular patient had to be moved to the isolation room for medical reasons since the time of the morning meeting, and this information was available via the patient information system. This explanation would not be useful if it had not been established already that we are on a ward round and the aim was to visit the patients. Vice versa, if the system would generate a justification for its assumption of being on a ward round, this would still not satisfy the need of the user to know why he should go to the isolation room.

The explanation shown to the user is pointing out the relevance of performing a particular action, namely visiting the isolation room.

6 Summary and Future Work

This paper has presented an approach to combine explanations and case-based reasoning in context awareness. It has been argued that explanations can be viewed from both a system centric and user centric perspective. It has further been described how explanations play a key role in ambient intelligent systems, where traditional interaction has been changed, and division of responsibility has been shifted from user initiated to mixed initiative.

The current implementation does support some of the goals seen from an user centric perspective; that is *elucidating* why it assumes that a given situation is occurring. The *transparency* goal is satisfied by CREEK's ability to graphically represent the case-match and the underlying semantic network representation, as well as a textual representation of the overall match strength and individual feature matches. A textual representation of explanations used in the reasoning process is also shown to the user.

The *conceptualization* goal is supported by providing a means to the user to explore the relation of the case to the underlying knowledge base. By examining the relations of the case features to this underlying knowledge structure, the user can gain some insight into the conceptual model the system has of the domain at hand.

The *justification* goal is to some degree supported in our current implementation. CREEK does explain why two different features are similar by means of plausible inheritance [28].

The most important future development is to apply the facets of explanation to the context sensitive part of our system. Here the Decomposer and Application agents must be able to use a system centric perspective on identifying a suitable behaviour in a given situation; as well as explicating why a given behaviour was chosen to the user.

Another direction for future research we want to explore is to tie our findings back in to our earlier work on a design methodology for explanation-aware intelligent systems with a socio-technical perspective. Our research has shown that ambient intelligent systems can benefit from a psychologically plausible knowledge model, but we have not yet explored the relation to the different knowledge containers in detail.

Additionally, we want to augment existing design guidelines with methods for the analysis of social aspects that can lead to a better understanding than ad-hoc methods. It is also important to note that we have not yet fully utilised some aspects of our theoretical foundations in organisational psychology, like the notions of breakdown situations or functional organs in Activity Theory, or the use of Semiotics for the organisation of the user interface itself.

References

1. Ducatel, K., Bogdanowicz, M., Scapolo, F., Leijten, J., Burgelman, J.C.: ISTAG scenarios for ambient intelligence in 2010. Technical report, IST Advisory Group (2001)
2. Marsh, S.: Exploring the socially adept agent. In: Proceedings of the First International Workshop on Decentralized Intelligent Multi Agent Systems (DIMAS 1995). (1995) 301–308
3. Pieters, W.: Free will and intelligent machines. Technical report, NTNU Trondheim (2001)
4. Schank, R.C.: Dynamic memory; a theory of reminding and learning in computers and people. Cambridge University Press (1982)
5. Schank, R.C.: Explanation Patterens – Understanding Mechanically and Creatively. Lawrence Erlbaum, New York (1986)
6. Sørmo, F., Cassens, J., Aamodt, A.: Explanation in case-based reasoning – perspectives and goals. Artificial Intelligence Review **24** (2005) 109–143
7. Roth-Berghofer, T.R.: Explanations and case-based reasoning: Foundational issues. In Funk, P., Calero, P.A.G., eds.: Proceedings of the 7th European Conference (ECCBR 2004). Volume 3155 of LNCS (LNAI)., Madrid, Spain, Springer (2004) 389–403
8. Richter, M.M.: The knowledge contained in similarity measures. Invited Talk at the First International Conference on Case-Based Reasoning (ICCBR 1995), Sesimbra, Portugal (1995)
9. Roth-Berghofer, T.R., Cassens, J.: Mapping goals and kinds of explanations to the knowledge containers of case-based reasoning systems. In Muñoz-Avila, H., Ricci, F., eds.: Case Based Reasoning Research and Development (ICCBR 2005). Volume 3630 of LNCS (LNAI)., Chicago, IL, Springer (2005) 451–464
10. Kofod-Petersen, A., Cassens, J.: Using activity theory to model context awareness. In Roth-Berghofer, T.R., Schulz, S., Leake, D.B., eds.: Modeling and Retrieval of Context: Second International Workshop, MRC 2005, Revised Selected Papers. Volume 3946 of LNCS (LNAI)., Edinburgh, UK, Springer (2006) 1–17
11. Weiser, M.: The computer for the 21st century. Scientific American (1991) 94–104
12. Díaz-Agudo, B., González-Calero, P.A.: An architecture for knowledge intensive CBR systems. In Blanzieri, E., Portinale, L., eds.: Advances in Case-Based Reasoning: 5th European Workshop, EWCBR 2000. Volume 1898 of LNCS (LNAI)., Trento, Italy, Springer (2000) 37–48
13. 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). Number 3155 in LNCS (LNAI), Madrid, Spain, Springer (2004) 1–15
14. Zimmermann, A.: Context-awareness in user modelling: Requirements analysis for a case-based reasoning application. [29] 718–732

15. 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 (2005) 112–115
16. Bénard, R., Bossard, C., Loor, P.D.: Context's modelling for participative simulation. In: Proceedings of the 19th FLAIRS Conference, AAAI Press (2006) 613–618
17. Kofod-Petersen, A., Aamodt, A.: Case-based situation assessment in a mobile context-aware system. In Krüger, A., Malaka, R., eds.: Artificial Intelligence in Mobile Systems (AIMS 2003), Universität des Saarlandes (2003) 41–49
18. 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 LNCS (LNAI)., Ölüdeniz, Turkey, Springer (2006) 211–225
19. Kofod-Petersen, A., Mikalsen, M.: Context: Representation and reasoning – representing and reasoning about context in a mobile environment. Revue d'Intelligence Artificielle **19** (2005) 479–498
20. Göker, A., Myrhaug, H.I.: User context and personalisation. In: Workshop proceedings for the 6^{th} European Conference on Case Based Reasoning. (2002)
21. Kofod-Petersen, A., Mikalsen, M.: An architecture supporting implementation of context-aware services. In Floréen, P., Lindén, G., Niklander, T., Raatikainen, K., eds.: Workshop on Context Awareness for Proactive Systems (CAPS 2005), Helsinki, Finland, HIIT Publications (2005) 31–42
22. 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 (2006) 619–624
23. Leake, D.B.: Workshop introduction: ECCBR-04 workshop on explanation in CBR. Slides, https://www.cs.tcd.ie/research_groups/mlg/ecbrws2004/, last access 22.02.2007 (2004)
24. Cunningham, P., Doyle, D., Loughrey, J.: An Evaluation of the Usefulness of Case-Based Reasoning Explanation. [29] 122–130
25. 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 (1991) University Microfilms PUB 92-08460.
26. Majchrzak, A., Gasser, L.: On using artificial intelligence to integrate the design of organizational and process change in US manufacturing. AI and Society **5** (1991) 321–338
27. Gregor, S., Benbasat, I.: Explanations from intelligent systems: Theoretical foundations and implications for practice. MIS Quarterly **23** (1999) 497–530
28. Sørmo, F.: Plausible inheritance semantic network inference for case-based reasoning. M.sc. thesis, Norwegian University of Science and Technology (NTNU), Department of Computer and Information Science, Trondheim (2000)
29. Ashley, K.D., Bridge, D.G., eds.: Case-Based Reasoning Research and Development: Proceedings ICCBR 2003. Number 2689 in LNCS (LNAI), Trondheim, Norway, Springer (2003)