Explanatory Capabilities in the CREEK Knowledge-Intensive Case-Based Reasoner

Anders KOFOD-PETERSEN [1], Jörg CASSENS, Agnar AAMODT
Department of Computer and Information Science (IDI), Norwegian University of Science and Technology (NTNU), Trondheim, Norway

Abstract. The ability to give explanations for its reasoning and behaviour is a core capability of an intelligent system. There are a number of different goals a user can have towards such explanations. This paper presents how the knowledge intensive case-based reasoning framework CREEK can support some of these different goals in an ambient intelligence setting.

Keywords. Case-based reasoning, ambient intelligence, explanation

Introduction

Explanations have been identified as one of the most important aspects of intelligent systems in general [1,2,3], and for ambient intelligent systems in particular [4]. When systems are assigned a kind of responsibility from their users, and exhibit pro-active behaviour, explanations are often the most important way to instil trust. This is especial true in ambient intelligent systems where the main interface often is behavioural.

Recent developments in ubiquitous and pervasive computing have shown that to achieve the visions proposed, systems must have far more complicated capabilities than initially identified by Weiser [5]. This has lead to the developments jointly labelled as *ambient intelligence* [6]. Ambient intelligence is defined as a system's ability to appreciate its environment, be aware of persons present, and respond to these persons needs in an intelligent manner.

We have earlier demonstrated how *case-based reasoning*, combined with a socio-technical analysis of the domain, can be utilised as a means of achieving ambient intelligence [7,8]. The use of case-based reasoning [9] is partly motivated by understanding reasoning as an explanation process [10]. Our understanding of similar occurrences of a situation assist us in comprehending stories, in such a way that details omitted or implicitly assumed do not make a story incomprehensible.

We have previously presented a framework for explanations in intelligent systems with a special focus on case-based reasoning [2]. Specifically, we identified five goals that explanations can satisfy.

[1] Corresponding Author: Department of Computer and Information Science, Norwegian University of Science and Technology, 7048 Trondheim, Norway; E-mail: anderpe@idi.ntnu.no

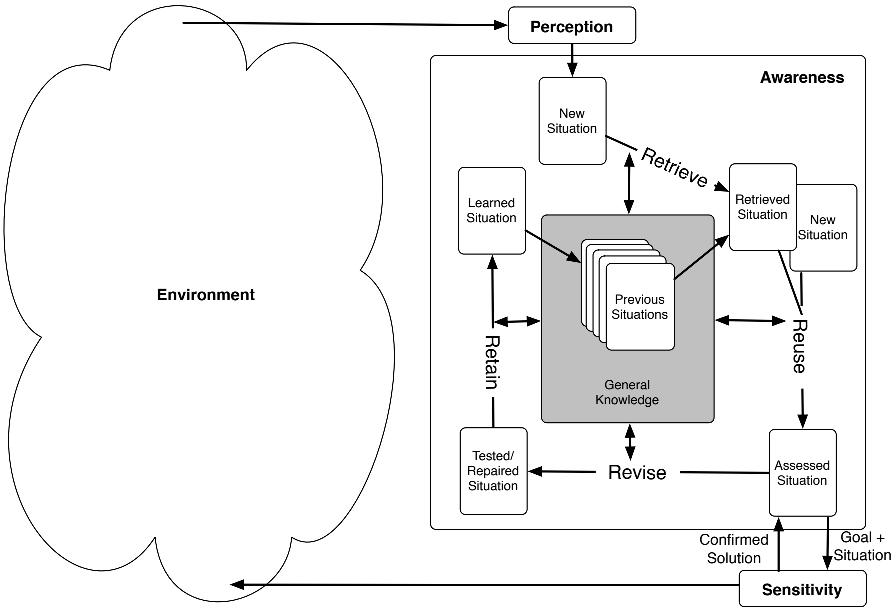

Figure 1. The Case-based Reasoning Cycle in Ambient Intelligence

The work presented here demonstrates the current state of the explanatory capabilities of the knowledge-intensive case-based reasoning system CREEK [11]. In particular viewed in the light of the five explanation goals identified in [2].

The rest of the paper is organised as follows: Firstly, a short introduction to the CREEK system, and how it is used in an ambient intelligent setting, is given. Secondly, an overview of the five explanation goals and how they fit into the ambient intelligence paradigm is given. This is followed by a description of how CREEK achieves the three goals which it is capable of supporting. Finally, a short summary and an outlook on future work is presented.

CREEK and Ambient Intelligence

Case-based reasoning approaches the reasoning process, not by the classic approach of rules or general knowledge, but by using episodic memory. The knowledge base consists of situations experienced by the system and the solution to the problem in these situations (cases).

A case-based reasoning system can, to some degree, be perceived as a rule-based system with very long rules. A case's findings can be viewed as a problem's antecedent and its solution as the consequence. However, they differ in several important ways: *i)* rules are patterns, whereas cases are constants; *ii)* rules are retrieved based on exact matching, where cases only require partial matching; *iii)* rules are small, independent, and consistent pieces of domain knowledge, cases can be large chunks of, potentially redundant, domain knowledge.

Table 1. Context and 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

Case-based reasoning has commonly approached the matching of a newly encountered situation to the existing case base by comparing *surface* features. The features are the findings, or description of a case, and are often represented as attribute-value pairs. Other approaches include comparison of *structural* properties. Structural features may contain explicit structural information, i.e. internal dependencies among features within the case itself, or implicitly through an additional model of general domain relationships that contains concepts referred to in the case. Such an approach tends to be more computational expensive than purely syntactical comparison, yet it often produces more relevant cases [12].

The CREEK method [11,13], which is used in the work presented here, is a *knowledge-intensive* case-based reasoning architecture especially developed to approach problem-solving and learning in open and weak-theory domains. CREEK contains structural features by linking flat case structure to a multi-relational network of concepts in the case. The main asset of CREEK is the fact that the cases are submerged into the general domain knowledge. This model is realised through a multi-relational semantic network, where object-oriented, frame-based representation is used to capture both cases and domain knowledge.

Modelling of knowledge in CREEK is a combination of a top-down process for the initial knowledge acquisition, and a bottom-up process of continuous learning by retaining cases. The top-down process is to acquire and develop the conceptual model required to define the domain model, to define the case contents and structure, and to manually described an initial set of case. To facilitate this knowledge acquisition, and the reasoning process, CREEK is equipped with a top-level ontology, which encompasses the high-level concepts and relations required.

CREEK has been experimentally used within an ambient intelligent environment for hospital wards [8]. Figure 1 depicts how the case-based reasoning cycle [9] has been adapted to an ambient intelligent environment. The knowledge model was developed by executing an ethnographical study at the local university hospital [7,14].

The Five Explanation Goals

Sørmo et al. [2] describe five different explanation goals that an intelligent system should be able to 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. *Conceptualisation* 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. These goals are defined from the perspective of a human user. His expectations on what constitutes a good explanation is situation dependent and has a historic dimension [15].

We have previously described how context awareness and sensitivity are related to explanations seen from either a user-centric or system-centric perspective [8]. Table 1 describes these relationships. In previous work we have described how all five goals are related to an ambient intelligent system. In the work presented here, we will only investigate the three goals that CREEK currently supports.

Explanation Goals in CREEK

Transparency Goal

The transparency goal is concerned with a system's ability to explain how an answer was reached. In the case of the user centric perspective, a system can elucidate (see Table 1) why it assumes that a particular situation has been identified correctly.

CREEK approaches this by visualising why a known situation (case) matches the ongoing situation. Figure 2 depicts the matching of two cases: `US_V_OL9_1302_Car` being the unsolved new case and `V_OL9_1305_Car_S` being the known case. CREEK displays the way the two cases are related, thus achieving the transparency goal.

If we look at the new case, we can see that by following the relations through to the know case we are aware of the match. `US_V_OL9_1302_Car` is connected to `US_UC_V_OL9_1302_Car` through the `has context` relation. The `has context` relation points to a context that encapsulates all findings of a case. This context has several parts, all connected through the `has part` relation. If we look at the `has part` relation that connects to `EC_V_OL9_1302_Car` (the environmental context), is again connected to the person `SPL8`, who is a nurse[2]. `SPL8` is again connected to the unsolved case through `EC_V_OL9_1305_Car_S` (environmental context) and `UC_V_OL9_1305_Car_S`.

By looking at Figure 2 we can also observe that both cases are connected through other findings. We can observe that both cases share the `GroupLeaderRole` and `ExaminationResponsibleRole`. Finally, both cases have a `Time` that is very close to the other: `1040` versus `1002`. We will examine the closeness of these times in the justification goal.

Justification Goal

The justification goal is closely related to the transparency goal. Where transparency is concerned with presenting the reasoning trace, justification deals with the ability to explain why an answer is good. Justification is often preferable over transparency, as simply displaying the reasoning trace is not always sufficient, it can even be counter productive [16,17]. In the case of the user-centric perspective, a system can elucidate (see Table 1), by justification, why it has classified a situation correctly.

CREEK allows the user to investigate why two concepts are matched, in particular when they do not match syntactically. When using the CREEK interface it is possible to investigate all the findings used when matching two cases. This is of particular interest when matches occur that are not based on surface features, or syntactical match, but rather on semantical similarity.

[2] The fact that SPL8 is a nurse can be observed by exploring the knowledge base.

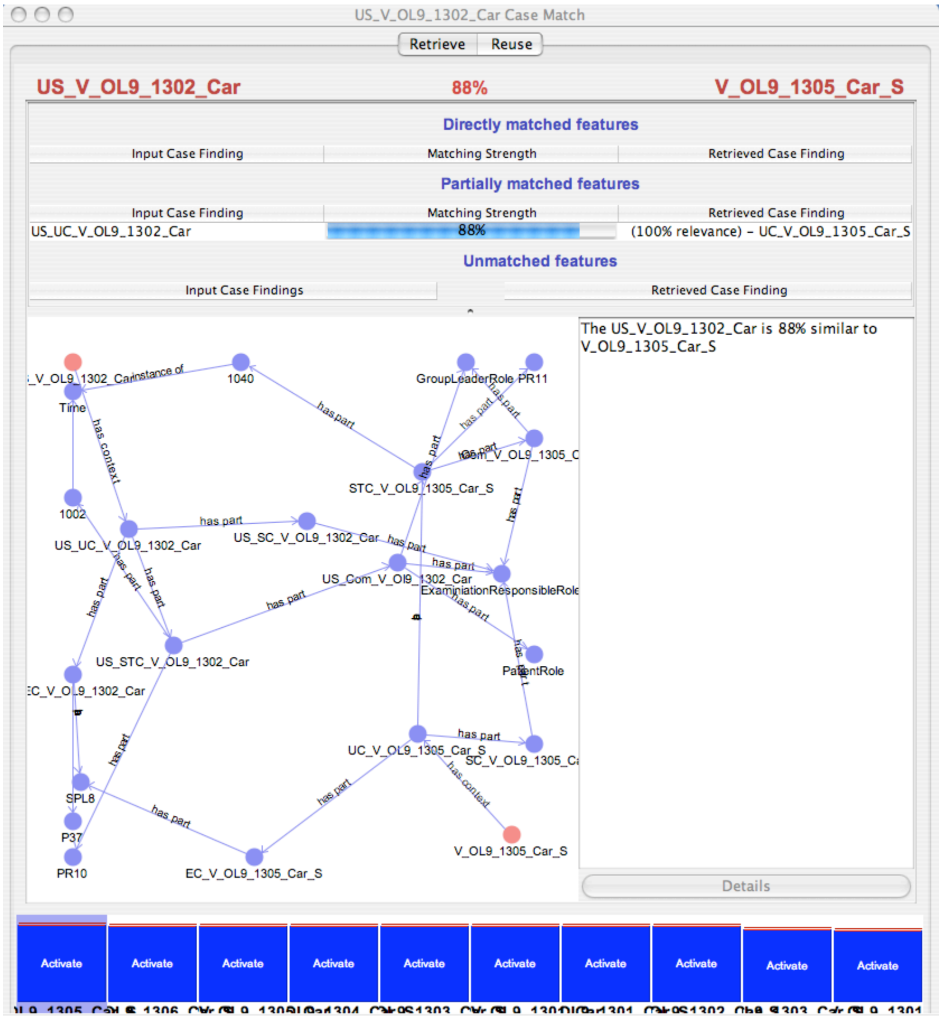

Figure 2. Transparency goal in CREEK

For example, if we look at the time in the two example cases, we can see that one of the situations is occurring at `1002` and one at `1040`. Looking at Figure 3 we can see that CREEK calculates that the two values are 95% similar, as times varies from 09:50 to 16:10[3]. In this case, CREEK has justified its assumption that the two points in time are close.

Conceptualisation Goal

The conceptualisation goal deals with the meaning of concepts. When we examined Figure 2 we found that both situations (cases) contained `SPL8`, and we noted that it was

[3]Please note that due to internal representational issues time is represented as integers. Thus, time is not correctly represented. However, for matching purposes these values are sufficient.

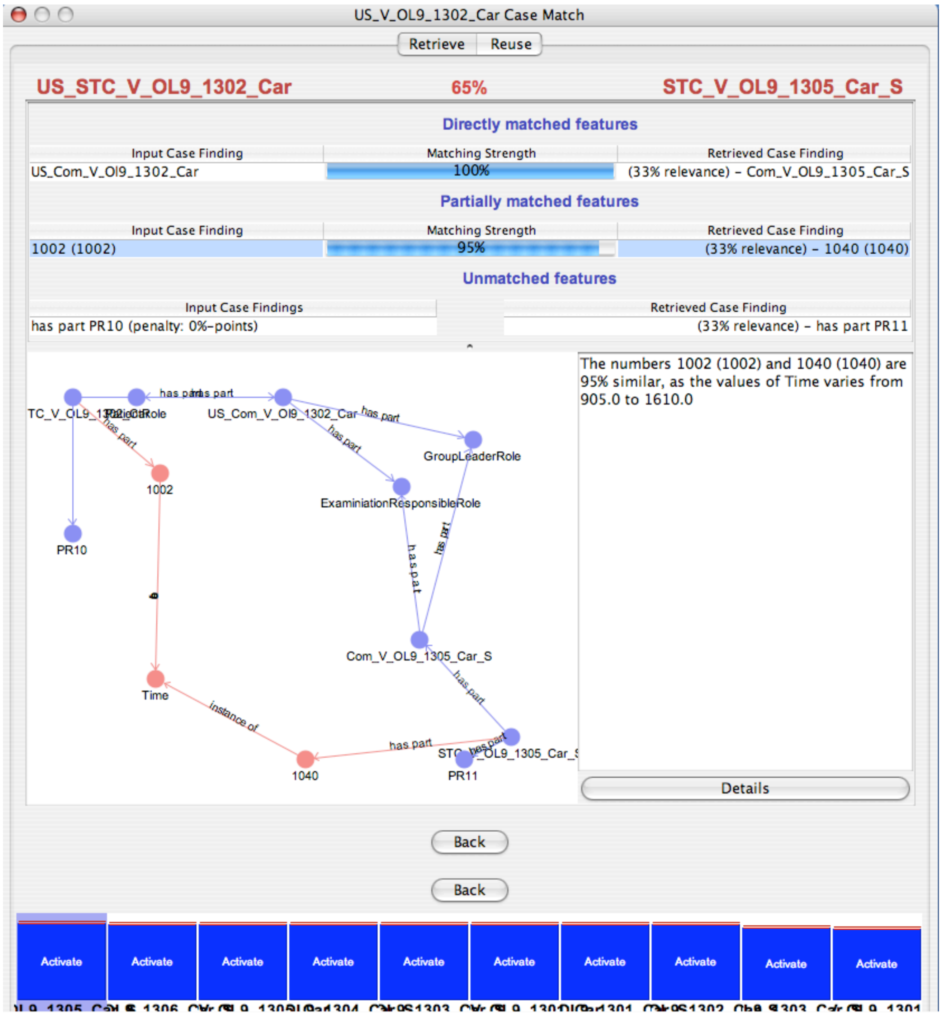

Figure 3. Justification goal in CREEK

a nurse. However, to explain the concept SPL8 we can examine the knowledge base contained in CREEK. Figure 4 shows a small subset of the knowledge base used.

When we examine the knowledge base, we can start with the concept Thing. According to the top-ontology in CREEK all concepts are at some point, through potential many levels of the multi-relational semantic network, either a subclass or an instance of Thing.

If we look closer at Nurse we can see that it Consumes an information source name FAMSOS, which is short for family and social issues. This information is offered by (among others) Patients. We can further examine the specific Nurse instance SPL8, and we can see that it is cast in the Role of GroupLeaderRole. Going back to Figure 2, this role was present in both cases. Finally, the location of SPL8 is LK4, which is an instance of a DoctorOffice, which again is a subclass of a Location.

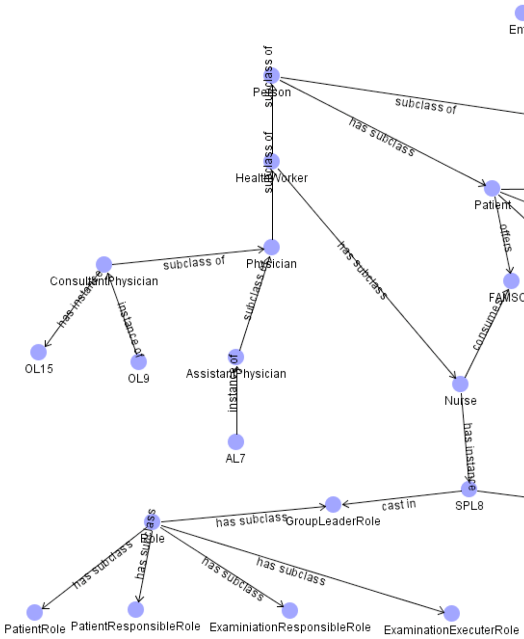

Figure 4. Conceptualisation goal in CREEK

Inspecting the knowledge model is the means of achieving the conceptualisation goal in CREEK.

Summary and Future Work

In this paper, we have described how the knowledge intensive case-based reasoner CREEK can support different user goals. We have taken a user-centric perspective, enhancing the system's communication with the user. In particular, we have shown examples for how we support the *transparency*, the *justification*, and the *conceptualisation* goals. For *transparency* and *justification*, we have focused on elucidating the system's reasoning, that means we have described how the system delivers an explanation to the user which explains why a particular situation was identified. Support of the *conceptualisation* goal in the way we have outlined can be useful both for elucidation and explication, which is an explanation targeting the behaviour of the system.

One of the future research areas we are currently exploring is the support of the *relevance* and the *learning* goal. The learning goal is special in the sense that it focuses on the user's interest in the application domain (hence the real world), and not on some particular behaviour of the system. It is mainly important to be supported in intelligent tutoring applications. Therefore, we concentrate on the *relevance* goal.

Another area of future work deals with the question of how to support the different user goals with the means of behavioural interfaces. The main form of interaction engaged in by the user of an ambient intelligent system is determined by the behaviour of that system. This implies that we have to focus on means to transport explanations on different channels than traditional computer displays.

Acknowledgements

Part of this work has been supported by Accenture Innovation Lab Norway.

References

[1] Leake, D.B.: Evaluating Explanations: A Content Theory. Lawrence Erlbaum Associates (1992)
[2] Sørmo, F., Cassens, J., Aamodt, A.: Explanation in case-based reasoning – perspectives and goals. Artificial Intelligence Review **24** (2005) 109–143
[3] 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 LNAI., Chicago, Springer (2005) 451–464
[4] Kofod-Petersen, A., Cassens, J.: Explanations and context in ambient intelligent systems. In kokinov, B., Richardson, D.C., Roth-Berhofer, T.R., Vieu, L., eds.: Modeling and Using Context, proceedings of the 6th International and Interdisciplinary Conference (CONTEXT 2007). Volume 4635 of Lecture Notes in Artificial Intelligence., Roskilde, Denmark, Springer Verlag (2007) 303–316
[5] Weiser, M.: The computer for the 21st century. Scientific American (1991) 94–104
[6] Ducatel, K., Bogdanowicz, M., Scapolo, F., Leijten, J., Burgelman, J.C.: ISTAG scenarios for ambient intelligence in 2010. Technical report, IST Advisory Group (2001)
[7] 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 Lecture Notes in Computer Science., Edinburgh, UK, Springer Verlag (2006) 1–17
[8] 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 (2006) 211–225
[9] Aamodt, A., Plaza, E.: Case-based reasoning: Foundational issues, methodological variations, and system approaches. AI Communications **7** (1994) 39–59
[10] Schank, R.C.: Explanation Patterens – Understanding Mechanically and Creatively. Lawrence Erlbaum, New York (1986)
[11] 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
[12] López de Mántaras, R., McSherry, D., Bridge, D., Leake, D., Smyth, B., Craw, S., Faltings, B., Maher, M.L., Cox, M., Forbus, K., Keane, M., Aamodt, A., Watson, I.: Retrieval, reuse, revision, and retention in case-based reasoning. Knowledge Engineering Review **20** (2005) 215–240
[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 (1991) University Microfilms PUB 92-08460.
[14] 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
[15] Leake, D.B.: Goal-based explanation evaluation. In: Goal-Driven Learning. MIT Press, Cambridge (1995) 251–285
[16] 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
[17] Gregor, S., Benbasat, I.: Explanations from intelligent systems: Theoretical foundations and implications for practice. MIS Quarterly **23** (1999) 497–530