Explanations and Context in Ambient Intelligent Systems

Anders Kofod-Petersen and Jörg Cassens

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

Abstract. Ambient intelligent systems are context aware by perceiving and reasoning about their environment, they perceive the needs of their users and proactively respond to these needs by being context sensitive. Users do not interact with these systems by traditional means only, but also through behavioural interfaces. This combination of mixed initiative systems and unconventional interfaces puts strong requirements on the explanatory capabilities of any system. The work presented here focuses on explaining the behaviour of an ambient intelligent systems to its users. It demonstrates how explanations can be combined with context to deal with the different types of explanations that are required for a meaningful interaction of a system and its users.

1 Introduction

Recent insights into both ubiquitous and pervasive computing have lead to the realisation that to achieve the scenarios and visions proposed, systems must be viewed as more complex than initially argued by Weiser [1]. This has lead to the developments jointly labelled as *ambient intelligence* [2]. The explicit focus on intelligence stands in stark contrast to the original argument by Weiser, where: "no revolution in artificial intelligence is needed – just the proper embedding of computers into the everyday world" [1, p. 3].

The core of an ambient intelligent systems is the ability to appreciate the system's environment, be aware of persons in this environment, and respond intelligently to their needs. To realise the abilities of an ambient intelligent system, three main areas of responsibility can be identified [3]: first, the initial responsibility of *perceiving* the world that the system inhabits; second, the responsibility of being aware of the environment and reason about ongoing situations, which traditionally has been labelled as *context awareness*; and third, exhibit appropriate behaviour in ongoing situations by being *context sensitive* [3,4].

Arguably one of the most important aspects of an ongoing situation is the activity that is occurring. For an ambient intelligent system to function it must be able to reason about its own, as well as other ongoing activities. Systems that aim at exhibiting the properties connected with ambient intelligence must be more than mere reactive systems, where deliberation and reasoning plays an important part.

Marx [5] demonstrates this difference by arguing that even though a spider conducts operations that resemble those of a weaver, and a bee humbles many an architect, there is a significant difference between them. Even the worst architects raise the structures in their imagination before they are erected in reality. At the end of each labour-process, we get a result that already existed in the imagination of the labourer. The labourer not only affects the materials used, but also realises a purpose that gives the law to his modus operandi, and to which the labourer's will must be subordinated. Besides the exertion of the bodily organs, the process demands that, during the whole operation, the workman will be steadily in consonance with his purpose. This means close attention. The less he is attracted by the nature of the work, and the mode in which it is carried on, and the less, therefore, he enjoys it as something which gives play to his bodily and mental powers, the more close his attention is forced to be.

The elementary factors of the labour-process are: *i*) the personal activity of the labourer, *ii*) the subject of the work, and *iii*) its instruments.

This is also the starting point for activity theory. To capture the activity related aspects of any situation, activity theory [6,7] can be used to acquire and model the relevant knowledge. Briefly, activity theory considers activities as a set of actions and operations on an object. These actions and operations are conducted by a labourer, or subject, to achieve an already imagined outcome. The subject's actions and operations are mediated by the use of certain instruments, or artefacts. We will elaborate on this later.

According to Turing [8], one indication that a system *is* intelligent is its ability to *appear* intelligent; i.e. by passing the Turing test. Therefore, we need to understand what makes humans appear intelligent.

Following Kant, human understanding has as a necessary constituent the ability to conceptualise perceived phenomena (structured through 'categories of understanding') through an active, discursive process of making sense of the intuitive perception [9, p. 58].

In later works, Kant gives us a more detailed description of his understanding of human reason. He makes clear that the human ability of reasoning has perceptivity (attentio), abstraction (abstractio), and reflection (reflexio) as its necessary preconditions [10, p. 138].

Further on, it is important to note that the ability of human beings to act freely, the ability to initiate a causal chain from freedom, is coupled with his ability to act morally (Kant describes freedom as the ratio essendi of the moral law, while the moral law is the ratio cognoscendi of freedom [11, p. 4]). Kant couples the ability to act morally (and thus freely) with the ability to give a rational explanation of the behaviour in his categorical imperative – "Act so that the maxim of thy will can always at the same time hold good as a principle of universal legislation" [11, p. 30]. Therefore, we can ascribe the ability of explaining ones behaviour and motives to every rational being, that means to every intelligent entity. We therefore count explanatory capabilities, in particular the ability to explain ones own understanding of the world and ones own behaviour, as a necessary precondition for appearing intelligent.

2 Background

One approach to realising intelligent behaviour is by employing *case-based reasoning* [12]. This method springs from understanding reasoning as an explanation process [13]. Our understanding of common occurrences assists us in comprehending stories, in such a way that details omitted or assumed implicitly do not make a story incomprehensible for us. Our general knowledge about situations, the expectations, and the behaviour which should be exhibited are stored in what has been referred to in psychology as *scripts* [14], which are closely related to the concept of *schema* [15,16].

Sørmo et al. [17] 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. These goals are defined from the perspective of a human user. His expectation on what constitutes a good explanation is situation dependend and has a historic dimension (compare e.g. Leake [18]).

Roth-Berghofer has explored some fundamental issues with different useful kinds of explanations and their connection to the different knowledge containers of a case-based reasoning system [19]. Based on earlier findings from natural language explanations in expert systems, 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 [20], namely case base, similarity measure, adaptation knowledge, and vocabulary.

Building on these two works, we have earlier started to investigate a combined framework of user goals and explanation kinds [21]. 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 further on 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 a socio-psychological analysis of workplaces in order to be able to design systems that can meaningful engage in socio-technical interactions.

In order to further explore the assumed advantages of designing systems from a socio-technical perspective, we have later on investigated the use of theories from industrial and organisational psychology in designing a case-based reasoning system geared towards ambient intelligence. The work presented here shows how the user-centric explanation goals can be satisfied by relating kinds of explanations in context awareness and context sensitivity with a socio-technical approach to modelling context.

3 Use of Activity Theory as a Means to Model Context

We have published some work on using activity theory to model context awareness [22,23]. Although we have discussed the use of this theoretical framework to help understand when to deliver an explanation [24], we have not previously explored how to make use of the same theoretical framework for designing explanatory capabilities for context aware systems. We will now outline how these deficiencies can be overcome.

First in this section, we will give a short summary of aspects of activity theory that are important for this work. See [25] for a short introduction to activity theory and [26,27] for deeper coverage. The theoretical foundations of activity theory in general can be found in the works of Vygotsky and Leont'ev [6,7,28].

Activity theory is a descriptive tool to help understand the unity of consciousness and activity. Its focus lies on individual and collective work practice. Some of its basic properties are:

- **Hierarchical structure of activity:** Activities (the topmost category) are composed of goal-directed actions. These actions are performed consciously. Actions, in turn, consist of non-conscious operations.
  If an action fails, the operations comprising the action can get conceptualised and might become conscious actions in the next attempt to reach the overall goal. This is referred to as a *breakdown situation.*
- **Object-orientedness:** Objective and socially or culturally defined properties. Our way of doing work is grounded in a praxis which is shared by our co-workers and determined by tradition. The way an artefact is used and the division of labour influences the design. Hence, artefacts pass on the specific praxis they are designed for.
- **Mediation:** Human activity is mediated by tools, language, etc. The artefacts as such are not the object of our activities, but appear already as socio-cultural entities.
- **Continuous Development:** Both the tools used and the activity itself are constantly reshaped. Tools reflects accumulated social knowledge, hence they transport social history back into the activity and to the user.
- **Distinction between internal and external activities:** Traditional cognitive psychology focuses on what is denoted internal activities in activity theory, but it is emphasised that these mental processes cannot be properly

understood when separated from external activities, that is the interaction with the outside world.

We have used an expanded model of activity theory, the Cultural Historical Activity Theory (CHAT, compare e.g. [29,30]), in order to analyse the use of technical artefacts as instruments for achieving a predefined goal in the work process as well as the role of social components, like the division of labour and community rules.

We have linked these different aspects of an activity to related categories of context in order to build a psychologically plausible context model. At the same time, we have used the model to guide our analysis of the work processes to be modeled into the system. However, we have not exploited all of the above mentioned aspects of activity theory in order to gain insight into the expectations and needs of the prospective users with regard to explanations.

4 Explanations and Context

The term explanation can have different foci. Either as goals that explanations can satisfy or as kinds of explanations that can be given. In addition, Leake identifies three different facets of explanations within the context of case-based reasoning [31]:

- 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. [17,32]). We are left with two functions of an explanation, as described in [33]: 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:[1]

- 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.

[1] This distinction is valid not only for case-based reasoning systems.

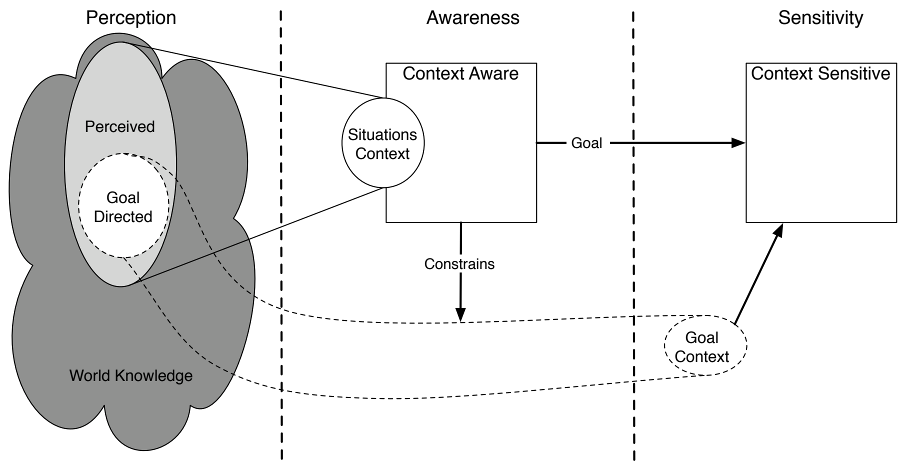

Fig. 1. Dual use of Context

We have earlier argued that an ambient intelligent system consists of three layers, each with their own responsibility [3]. The top layer is responsible for *perceiving* the world and order the perceived data into a coherent context structure on which reasoning is possible. The awareness layer is responsible for assessing the context and classify an ongoing situation. This layer demonstrates the ability of *context awareness*. Finally, the third layer is responsible for selecting and executing suitable behaviour based on the classification done in the awareness layer. This ability is referred to as *context sensitivity*.

In this architecture, context serves two purposes. Initially it is used as a focussing lens on the part of the world that can be perceived. Here the context limits the parts of the knowledge that the system uses to classify the situation. The second use of context is in the context sensitivity layer, where context is viewed as a lens that focuses the part of the system's knowledge that is to be used to satisfy the goal of the situation.

Figure 1 depicts the dual use of context. Initially the Situation Context is what the context aware part uses to execute the case-based reasoning process that classifies the situation. Once the situation has been classified a suitable goal for this situation is found. This goal further limits the part of the context that is necessary for the context sensitivity part to exhibit appropriate behaviour. The goal as well as the context are made available to the context sensitivity part, as is indicated by the Goal arrow and the Goal Context in the figure.

For the purpose of this paper we will disregard the perception layer of the architecture as the perception layer demonstrates no reasoning capabilities, and only structures perceived data syntactically. Following our earlier arguments introduced in [3], we identify these two aspects as two distinct steps in the reasoning process:

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

- **Context Awareness:** Trying to detect which situation the system is in.
  Example: An ambient intelligent system for supporting health personnel figures out 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.

Combining these views on explanation and on context, we end up with two dimensions of inquiry as depicted in Table 1. This table shows the four different areas where explanations can be required, divided into a system centric and user centric view.

In the system centric view where explanations are a part of the reasoning process it is possible to initially generate an explanation used to *recognise* the situation. In this step we are using explanations to find out what situation we are in, by explaining similarities between a new situation and previously experienced situations. Following the recognition of a situation we can now use explanations to identify appropriate behaviour.

When dealing with the user centric view we can initially use explanations to elucidate why the system assumes that we are in a certain situation. The system can use all available sources of knowledge in order to gain the user's confidence in its capabilities. In the situation where the the system is required to explicate the behaviour that it exhibits, the explanation is used to explain why it takes a specific action.

As described above, activity theory has been used to recognise contextual facets of a work situation. By integrating the knowledge necessary for supporting the different explanatory goals of the user with this contextual information, the explanatory capabilities of the system are coupled with the different contexts. Hence, the hypothesis is that this explanatory knowledge will indeed primarily be used in the appropriate context.

We will now explore the relations between the basic properties of activity theory and explanation goals.

Hierarchical structure of activity: The fact that activities are hierarchically structured, and that changes in these structures occur, facilitates certain explanation goals. Actions that are performed often will be transformed into operations. Vice versa, if an anticipated outcome of an operation does not occur, non-conscious operation will become conscious actions. This is called a breakdown situation. The explanatory capabilities of a system should support this. In fact two goals are relevant in these situations:

- **Transparency:** If parts of the non-conscious operations are carried out by artefacts, the system might need sufficient knowledge to explain the artefacts inner working in case of a breakdown.
- **Relevance:** If an artefact involved in an action can behave differently than expected, it should be made clear why the unexpected behaviour occurred.

Object-orientedness: In the activity theoretical sense of the term object-oriented, the meaning of this term is twofold. On one hand, it highlights that all human activities have an objective, a goal, and therefore points towards the mental part of an activity. On the other hand, it refers to the fact that this mental objectives are directed towards the physical world. This holds for automated processes insofar as the automation already assumes a goal, and is supposed to support this goal:

- **Transparency:** It should be possible for a system to explain its relation to the physical processes.
- **Justification:** An intelligent system should be able to explain its goals to the user.

Mediation: Every activity will incorporate some tools, be it physical (like machinery) or psychological artefacts (like language). If parts of the activity are carried out by an intelligent artefact, this artefact both acts as a mediator in the physical world and as a mediator of the psychological processes of the user:

- **Justification:** The system should be able to explain the connection between its actions and the reasoning process.

Continuous development: The aspect of continuous development deals with the continuous change in the way we interact with the world. Both the user's activities and the artefacts used are changing. It should be noted that this includes the necessity for an intelligent system to adopt to changes over time:

- **Learning:** The system should be able to support the user's learning processes. If the system is extended, or new capabilities are included, the system should be able to act as teacher. It should therefore incorporate knowledge about how the new component facilitates the problem solving process.

Distinction between internal and external activities: Activity theory tries to overcome the dichotomy of mental processes and the outside world by focussing on the relation between internal and external activities. It is therefore crucial that the system supports the user in building an understanding of the artefacts used.

- **Conceptualisation:** The system should support the user's understanding of it by providing means of explaining its own world model to him.

Not all explanation goals can be satisfied by an activity theoretical perspective alone. Some goals can only be satisfied by inspecting other parts of the knowledge model, either in all cases or for certain situations. As an example, when recognising a situation the transparency goal can be satisfied by supplying a trace of the reasoning process used for classification. The different sources of knowledge required to satisfy the different goals will be further discussed in the following section.

5 Identifying Explanations Kinds from Goals

With the combination of explanations and context described above it is now possible to identify different kinds of explanations by identifying the explanation goals of the user. The four different areas where explanations can be required are shown in Table 1. For the purpose of this paper we will focus on the user centric perspective where explanations are used to elucidate why the system identifies a particular situation and to explicate why a certain behaviour was chosen.

As we have stated before, users do not interact with an ambient intelligent system by traditional means only but also through behavioural interfaces. This means also that if the system gets everything right, it should be unobtrusive and supportive. The main situations where explanations are necessary are when something goes wrong, i.e. the system does not recognise the correct context or follows a path of actions which the user perceives as wrong, unusual, or unexpected. So while we do not dismiss the option that a user wants some explanation from the system even if it does what the user expects, we do not focus on this aspect in this paper. But it has to be kept in mind that the system's explanatory capabilities should also cover its ability to explain itself when nothing goes wrong. This is of special importance during the beginning of the use of the system in order to gain the user's trust into the system.

We would also like to point out that we have chosen not to consider the learning goal in this paper. The learning goal is specifically targeted towards educational systems. The goal of such a system is typically not only to find a good solution to a problem, but to explain the solution process to the user in a way that will increase his understanding of the domain. We do not consider this type of systems at the time being.

5.1 Context Awareness

In case of the context aware user centric perspective the system might *misclassify* a situation. In this case the system must satisfy the goals of *transparency, justification* and *conceptualisation*. In case of *transparency* and *justification* they explain the process through which the classification was reached. The choice between a *transparency* or *justification* goal is governed by the user's proficiency level. Where an expert user will require *transparency*, a novice user requires *justification*[2]. These two goals map to the 'how' and 'why' explanation kinds, where

[2] This separation will be used consistently throughout the rest of the paper.

'how' explains the causal chain of events leading to the classification, and 'why' justifies why the system thinks that the answer is good. The knowledge required to supply these kinds of explanations is found within the reasoning method, e.g. similarity measures in case-based reasoning.

5.2 Context Sensitivity

When dealing with the context sensitive user centric perspective, the system has two main situations in which explicating is required (not counting the situation where the system exhibits flawless behaviour). These two main situations are when the system exhibits *wrong behaviour* for the situation, and when it exhibits *unexpected behaviour*. Both of these situations can result in a breakdown situation as defined by activity theory. In case of a *wrong behaviour* the system's goal is not in line with the user's goal, any operations performed by the user will fail and become actions, thus a breakdown situation is occurring. In this case, the system displaying wrong behaviour must satisfy the same goals as when misclassifying the situation. This means that the *transparency/justification* goals must be satisfied. As with the case of misclassification, these goals map to the 'why' and 'how' kinds, which require knowledge about the reasoning process employed. In addition, the hierarchical structure principle in activity theory can guide the process through which these goals are satisfied.

From a user perspective, the system can *behave unexpectedly* in several different manners: it can request an *unexpected action* from the users, non-user actions can be performed by a *new or alternative person* or by a *new or unexpected artefact*.

When the system requests a *new action* from the user, a breakdown situation occurs, and the user must respond consciously. Again, the goals of *transparency/justification* must be satisfied. In addition, the system must satisfy the *relevance* goal by explaining the relevance of the requested action, and in case of previously unperformed actions *conceptualization* is required. The hierarchical structure principle in activity theory can guide the process through which the *transparency/justification* and *relevance* goals are satisfied, whereas the *conceptualization* goal can be satisfied by inspecting the specific domain model. The *relevance* goal maps to the 'purpose' kind where an explanation of the purpose of the requested action gives the relevance. Finally the *conceptualization* goal maps to the 'conceptualization' kind, mapping unknown concepts to known ones.

If a non-user action is performed by a *new or alternative artefact*, three goals must be satisfied: *transparency/justification, relevance* and *conceptualization*. The mediation principle in activity theory can guide the process where the *transparency/justification* goal is satisfied, whereas the *relevance* and *conceptualization* goals are satisfied by inspecting the specific domain model. As aforementioned the *transparency/justification* goals map to the 'why' and 'how' kinds and the *conceptualization* goal maps to the 'conceptualization' kind. In the case of the *relevance* goal, this maps to the 'purpose' kind when dealing with alternative artefacts by describing the purpose of this artefact. When dealing with a new artefact, the *relevance* goal also maps to the 'conceptualization' kind.

For non-user actions performed by a *new or alternative person*, the description is similar to the one for artefacts. However, one important distinction exists. In the activity theoretical part of the knowledge model persons are part of the community that cooperates with the user through a division of labour. However, our current modelling of persons and roles does not distinguish between the two. Thus, even though an action is performed by a new or alternative person, the fact that the role is unchanged means that the activity as viewed from the system is unchanged. This is contrary to the way artefacts are perceived, where using a new or alternative artefact to perform an action will result in a change in the activity. This means that no activity theoretical principle can guide the satisfaction of the goals, thus other parts of the knowledge model must be inspected.

6 Example

We will briefly investigate the relations between context and explanation by the means of an example. Let us consider the following scenario: We have a case-based diagnostic system for aircraft failures. An engineer is equipped with an intelligent mobile assistant and one of his tasks is to diagnose the probable causes of engine problems. Let us assume that the engineer is working both at his home base and at line stations where faults have occurred.

Scenario 1 – Misclassified Context: Let us assume that our engineer is going to work with the head of engineering on a new schedule for sending engineers to line stations. He is doing administrative work and not working on technical problems. The time of this meeting, however, is at a time where there is usually a briefing with all engineers, and the system also recognises that some of the other people usually participating at this meeting are present. However, instead of being in a meeting room, we are at the office of the head of the engineering group, a fact which contradicts the assumption of being in the briefing. The system might now explain away this unusual facet by generalising that both the meeting room and the office are rooms and that an office to a certain degree serves the same purpose as a meeting room. Therefore, the system assumes that we are in a briefing and delivers fault information about the airplanes which are scheduled to be worked on.

When this error becomes obvious to the engineer, he might want to check why the system displayed this kind of information. So we are in the *explicate* phase of Table 1. If he is an expert user of the system, he might have an interest what lead to the problem, so his goal is *transparency.* The kind of explanation helpful is a 'why' explanation, in particular one where the system displays the best matched cases and that it has classified the office as a general kind of room.

Scenario 2 – New Artefact Used: Let us now assume that the engineer is working on a diagnostic task and, in the course of this task, needs access to some performance data. This is recognised by the system. The knowledge source usually used for this kind of data is temporarily not available, so the system queries a different system which was added recently. This comes as a suprise to the engineer who was not aware of neither the unavailability of the first system nor the existence of the second.

The engineer now wants to know why the data from the new system is helpfull, he has a *relevance* goal. This can be supported by a 'purpose' kind of explanation, and by inspecting its own domain knowledge, the system can describe the purpose of the new data source, for example by explaining that the new data source is a backup system for performance data.

7 Summary and Future Work

This paper builds on a view of ambient intelligence encompassing first an understanding of the situation (context awareness) and then decisions on behaviour (context sensitivity). It has been argued that in both phases, explanations can be viewed from a system centric as well as a user centric perspective. It has further been described how explanations play a key role in ambient intelligent systems as a necessary prerequisite for a system being perceived as intelligent by human users.

The conceptual framework presented here describes how explanations can be used in the different parts of an ambient intelligent system. Further on, it describes how knowledge about requirements for explanations which can fulfil different user goals can be gained. We have introduced a means of taking user goals into account which is both psychologically plausible and in line with the tradition in context aware computing.

We have further on outlined how an understanding for user goals can be obtained both from an activity theoretic analysis of the activity environment and from the general and domain specific knowledge encompassed in the system at hand.

We have described how different user goals for explanations are related to different kinds of explanation and by this have outlined what knowledge a system has to contain in order to fulfil the user's goals. However, we have not yet tied this into a detailed methodology for intelligent systems design.

The three layered conceptual architecture (perception – awareness – sensitivity) combined with our conceptual model of explanations in ambient intelligence gives a foundation for the development of explanation aware applications. The different goals a user might have towards explanations together with their mapping to kinds and the inclusion of socio-technical analytic methods help us integrating the explanatory capabilities of the application at an early stage of the design process.

Our current implementation of an ambient intelligent case-based reasoning system can cater to the system centric perspective of explanations to some degree, but this has to be developed further. Regarding the user centric perspective, the current application does support the transparency, conceptualization and justification goals, where the latter is only supported partially due to the underlying issues with plausible inheritance in the current Java implementation. For the other goals, further implementation work is necessary.

Another aspect that deserves further attention is our model of the division of labour. In order to reconcile our view on artifacts and humans, we have to find ways to integrate the modelling of different persons as subjects into our generic context model.

Additionally, we want to augment existing design guidelines with methods for the analysis of social aspects which can lead to a better understanding of the environment in which the ambient intelligent system has to function than ad-hoc methods can give. 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. Weiser, M.: The computer for the 21st century. Scientific American (1991) 94–104
2. Ducatel, K., Bogdanowicz, M., Scapolo, F., Leijten, J., Burgelman, J.C.: ISTAG scenarios for ambient intelligence in 2010. Technical report, IST Advisory Group (2001)
3. 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
4. Yau, S.S., Wang, Y., Karim, F.: Development of situation-aware application software for ubiquitous computing environments. In: 29th Annual International Computer Software and Applications Conference (COMPSAC'02). Volume 1., Edinburgh, Scotland, IEEE Computer Society (2002) 107–112
5. Marx, K.: Das Kapital, Band 1: Der Produktionsprozeß des Kapitals. Volume 23 of Marx Engels Werke (MEW). Dietz-Verlag, Berlin (1867)
6. Vygotski, L.S.: Mind in Society. Harvard University Press, Cambridge, MA (1978)
7. Leont'ev, A.N.: Activity, Consciousness, and Personality. Prentice-Hall (1978)
8. Turing, A.: Computing machinery and intelligence. Mind **LIX** (1950) 433–460
9. Kant, I.: Kritik der reinen Vernunft (2. Auflage). Akademie (1787)
10. Kant, I.: Anthropologie in pragmatischer Hinsicht. Akademie (1798)
11. Kant, I.: Kritik der praktischen Vernunft. Akademie (1788)
12. Aamodt, A., Plaza, E.: Case-based reasoning: Foundational issues, methodological variations, and system approaches. AI Communications **7** (1994) 39–59
13. Schank, R.C.: Explanation Patterens – Understanding Mechanically and Creatively. Lawrence Erlbaum, New York (1986)
14. Schank, R., Abelson, R.: Scripts, Plans, Goals and Understanding. Erlbaum (1977)
15. Bartlett, F.C.: Remembering: A study in Experimental and Social Psychology. Cambridge University Press (1932)
16. Rumelhart, D.E.: Schemata: The building blocks of cognition. In Spiro, R.J., Bruve, B., Brewer, W.F., eds.: Theoretical Issues in Reading and Comprehension. Lawrence Erlbaum Assoc Inc (1980)
17. Sørmo, F., Cassens, J., Aamodt, A.: Explanation in case-based reasoning – perspectives and goals. Artificial Intelligence Review **24** (2005) 109–143
18. Leake, D.B.: Goal-based explanation evaluation. In: Goal-Driven Learning. MIT Press, Cambridge (1995) 251–285

19. 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 Lecture Notes in Computer Science., Madrid, Spain, Springer Verlag (2004) 389–403
20. Richter, M.M.: The knowledge contained in similarity measures. Invited Talk at the First International Conference on Case-Based Reasoning, ICCBR'95, Sesimbra, Portugal (1995)
21. 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 Lecture Notes in Computer Science., Chicago, USA, Springer (2005) 451–464
22. 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
23. Cassens, J., Kofod-Petersen, A.: Using activity theory to model context awareness: a qualitative case study. In Sutcliffe, G.C.J., Goebel, R.G., eds.: Proceedings of the Nineteenth International Florida Artificial Intelligence Research Society Conference, Melbourne Beach, AAAI Press (2006) 619–624
24. Cassens, J.: Knowing what to explain and when. In Gervás, P., Gupta, K.M., eds.: Proceedings of the ECCBR 2004 Workshops. Number 142-04 in Technical Report of the Departamento de Sistemas Informáticos y Programación, Universidad Complutense de Madrid, Madrid (2004) 97–104
25. Nardi, B.A.: A brief introduction to activity theory. KI – Künstliche Intelligenz (2003) 35–36
26. Bødker, S.: Activity theory as a challenge to systems design. In Nissen, H.E., Klein, H., Hirschheim, R., eds.: Information Systems Research: Contemporary Approaches and Emergent Traditions. North Holland (1991) 551–564
27. Nardi, B.A., ed.: Context and Consciousness. MIT Press, Cambridge, MA (1996)
28. Vygotski, L.S.: Ausgewählte Schriften Bd. 1: Arbeiten zu theoretischen und methodologischen Problemen der Psychologie. Pahl-Rugenstein, Köln (1985)
29. Kutti, K.: Activity theory as a potential framework for human-computer interaction research. [27] 17–44
30. Mwanza, D.: Mind the gap: Activity theory and design. Technical Report KMI-TR-95, Knowledge Media Institute, The Open University, Milton Keynes (2000)
31. 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)
32. Cunningham, P., Doyle, D., Loughrey, J.: An evaluation of the usefulness of case-based reasoning explanation. In Ashley, K.D., Bridge, D.G., eds.: Case-Based Reasoning Research and Development: Proceedings ICCBR 2003. Number 2689 in Lecture Notes in Computer Science, Trondheim, Norway, Springer (2003) 122–130
33. 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.