A context model for knowledge-intensive case-based reasoning

Pinar Öztürk and Agnar Aamodt

Department of Computer and Information Science
Norwegian University of Science and Technology
N-7034 Trondheim, Norway
{pinar, agnar}@idi.ntnu.no

Abstract: Decision-support systems that help solving problems in open and weak theory domains, i.e. hard problems, need improved methods to ground their models in real world situations. Models that attempt to capture domain knowledge in terms of, e.g. rules or deeper relational networks, tend either to become too abstract to be efficient, or too brittle to handle new problems. In our research we study how the incorporation of case-specific, episodic, knowledge enables such systems to become more robust and to adapt to a changing environment by continuously retaining new problem solving cases as they occur during normal system operation. The research reported in this paper describes an extension that incorporates additional knowledge of the problem solving context into the architecture. The components of this context model is described, and related to the roles the components play in an abductive diagnostic process. Background studies are summarized, the context model is explained, and an example shows its integration into an existing knowledge-intensive CBR system.

1 Introduction

The motivation behind this research is to improve artificial intelligence methods for problem solving and learning in *open and weak-theory domains* by developing an explicit model of context. Our example domain is medical diagnosis and treatment. Since there exists no neat theory from which to deduce conclusions for such domains, we have to rely on abductive methods that combine several types of available knowledge in interpreting a problem situation, generating and evaluating hypotheses, building explanations to support or reject them, etc. We are exploring the combined utilization of specific experiences in terms of *past cases* with a presumably extensive, multi-relational model of *general domain knowledge* (Aamodt, 1991). The purpose of introducing context as an explicit notion is partly to better 'ground' the knowledge model in the real world environment, i.e. to improve *solution quality*, and partly to more quickly focus on the right relations and concept types, i.e. to improve *problem solving efficiency.*

The *diagnostic process* involves two main subprocesses; generation of hypotheses and testing of these hypotheses. Hypothesis generation is an abductive process. Abductive inference allows *inference to the most plausible hypothesis* from a set of observations. Plausibility identifies what is rational, what reasonable people would agree on. A plausible hypothesis can later be retracted if new, conflicting information is provided in the subsequent process of hypothesis testing. Abductive inference is frequently used in expert-level reasoning and is consistently stated to be practiced by scientists, detectives, and physicians. It is ubiquitous in daily life as well. For example, upon seeing your neighbor walking into his house wet you may immediately think he has been swimming. However, swimming is not the only possible reason for one being wet. That is, your conclusion involving swimming is not guaranteed to be true. Other reasons could be that he walked home in the rain without an umbrella, that somebody poured water on him from the window of a building, or that children sprayed him with a hose. People reason quickly, even when there exist several alternative conclusions. They do not perhaps even need enumerating and considering all possibilities in order to come up with one. You may, for example, immediately adopt the "he was swimming" explanation if there is a swimming pool in the backyard, and he was wearing a swimming suit.

Computational models of abductive reasoning suffer from inefficiency. Investigating possible context utilization in abductive tasks such as diagnostic hypothesis generation should be able to improve abductive efficiency by focusing on only relevant regions of the explanation space. It has been acknowledged, in artificial intelligence as well as in cognitive science more generally, that context facilitates a selective and focused processing of information. Disagreements exist as to what context actually is, and in what way it affects the processes of problem solving and learning. An earlier trend was to consider context as a wholistic phenomenon, and to study context effects on various processes in a uniform way. This entails referring to context without making a distinction between the various types of contexts. Yet, different types of context elements have been observed to have significantly different effects on memory. Thus, the issue of the *role of context* is closely related to the issue of the distinction among several types of *context elements*.

In order to discuss relevant aspects of context as a kind of knowledge, a framework for knowledge analysis and modeling is needed. Recent research in the area of knowledge modeling and engineering has produced several methodologies for analyzing and modeling knowledge and information at a conceptual and implementation-independent level. This level of system description is often referred to as the *knowledge level*, after Newell's influential paper (Newell, 1982), while the level of implementation and representational constructs is referred to as the symbol level. Influential examples of methodologies that support knowledge analysis and modeling at the knowledge level are the Generic Tasks approach (Chandrasekaran, 1992), the COMMET methodology (Steels, 1992), and CommonKADS (Breuker, Van de Welde, 1994). So far, context has hardly been systematically studied within a knowledge level framework. By building on results from each of the three example methodologies, we have found this level of analysis useful in order to understand context as knowledge type, and what its links and dependencies to other knowledge types are. In particular, it allows to describe context in terms of *knowledge content*, and to take the first steps towards a content theory of context. As a result, we discuss context from two distinct but closely related perspectives, i.e. as a knowledge type, leading to the establishing of a *context ontology*, and as a means for shifting the focus of attention, which leads us to the idea of goal driven, *task-oriented shift of focus* during reasoning.

The paper describes a model of context for problem solving in weak-theory domains. This research is part of our general agenda, which is investigating the integration of case-specific (episodic) and general domain knowledge for solving weak-theory and open-domain problems. The paper is structured as follows: A rather thorough study of literature on context studies from AI-related parts of cognitive science has been undertaken. Essential findings are summarized in section 2.The following section (3) defines the main components of the context model. Section 4 describes the use of context in medical diagnosis. In section 5 our context model is incorporated in to the Creek architecture for knowledge-intensive case-based reasoning, while section 6 summarizes the paper.

2 Context studies in cognitive psychology

In investigating a content theory of context in problem solving, we try to exploit the interconnection between disparate disciplines, especially between cognitive psychology and AI. These two fields are intimately related. As Fisher & Yoo (1993) elegantly summarizes, "The common wisdom is that psychology provides a specification of intelligence, and AI provides tools for implementing this specification."

Context is studied differently in cognitive psychology than it would possibly be done in AI. Nevertheless, AI may utilize models offered by cognitive psychology. An important difference is that psychology is concerned with local tasks, whereas AI is concerned with cognition as a whole (Haugeland, 1984). Over the years, cognitive psychology has mainly focused on studying contextual effects related to relatively simple tasks, such as verbal learning and face recognition, rather than complex tasks, such as problem solving. Regarding the question of how these experiments can contribute to doing better AI, Polson (1984) writes that psychology can provide guidance in the identification of useful components for AI systems. That is exactly the way we utilize psychological studies of context in this work. For a broader discussion of this topic, see Özturk & Aamodt (1997).

Much empirical research has been devoted to investigating context. However, the focus of much of this research has not been on understanding the phenomenon itself, but rather to use context as a tool for investigating other topics (Wickens, 1987). An example is the early studies of memory retrieval aimed to study the relationships between the retrieval types of recall and recognition. The studies suggested context effects on retrieval, and context has been one of the primary mechanisms used to investigate memory theories. (Anderson and Bower, 1973; Kintsch, 1974; Thomson & Davies 1988). Continued studies of the relation between recall and recognition halerted some influential researchers to the importance of context as such (Tulving and Thompson, 1973; Eich, 1980; Baddeley (1982a and 1982b); Smith, 1986; Thompson and Davies, 1988). A technique that is commonly used for studying the influence of context on learning has been that the subjects first learn in one context and then are asked to remember in a context that differ in various ways from the learning context. Experiments on “state dependent learning”, environmental context effects, and semantic context effects in verbal learning and face recognition are the most famous ones. For example, in one study learning happened on land and was tested under water (Godden and Baddeley, 1975; Godden and Baddeley, 1980), in another, learning and remembering were tested in different rooms (see (Smith, 1986) for review). The results of these experiments agreed on the existence of context effects in retrieval but differed in their account of how context influenced recall and recognition as some did not find context effects in recognition while others did.

A dependency between the degree of context influence and the type of context utilized in a certain process has also been noticed. The attention is drawn to the variety in the types of contexts and to the variety in the ways that context may influence memory. These tangled findings triggered theoretical studies of memory which dealt with the characteristics of various kinds of context manipulations and with how and why context affects performance, rather than solely attempting to detect the context effects.

We will briefly mention some of these theories which have had major influence on our work. The *encoding specificity theory* maintains that, “An encounter with a word results in the creation of a unique trace....” (Watkins and Tulving, 1975), and implies that a subject does not learn, for example, words but *experiences*. The theory is based on the idea of a distinction between *episodic* and *semantic* memories. Encoding specificity theory, consequently, asserts that a better retrieval is conditioned on the closer resemblance between learning and retrieval situations, and thus makes clear the role of context elements both in learning and retrieval. What improves the retrieval is the degree of match (the similarity) between these two. Thus, the context in which encoding and retrieval take place plays a central role in determining what is remembered.

Craik and Lockhart’s (1972) *levels of processing hypothesis* explains how the nature of the process affects encoding. It elaborates the encoding specificity theory in that it notices that the depth of the process underlying learning determines what and how to encode (Craik and Tulving, 1975).

Baddeley suggested a distinction between *interactive* and *independent* context types. Baddeley's formulation is based on the way context influences learning, as well as retrieval. He argued that interactive context "influences the memory trace directly by affecting the way in which the stimulus is encoded", whereas independent context and stimulus are processed without interaction between them. Independent context bears a purely arbitrary relationship with the material learned. As such, it does not determine the interpretation of the material. Baddeley emphasizes that "The concept of interactive context places the emphasis on the *processing* carried out by the subject rather than on the characteristics of *stimulus material*" (Baddeley, 1982a).

This distinction between 'process' and 'characteristic of stimulus' has contributed to fostering our distinction between 'internal' and 'external' context types as the top level elements in our context ontology (section 3.2).

3 Towards a content theory of context

An ongoing debate is whether context is related to 'the state of the mind' or to 'the situation' (see (Dojat and Pachet, 1995) and (Kokinov, 1995)). We take a stance in the middle, since there is ample evidence that it has to do with both. The term 'situation' captures the 'ground facts' existing in a situation while 'situational context' comprises the contextual aspects pertinent to that situation. These are the facts that exist independently from the reasoner, in the sense that they are there before and after the reasoner notices them. The 'state of the mind' component of a context emerges while the reasoner solves a problem, and captures the goals, interests, and information needs of the reasoner in the progress of problem solving.

3.1 The role and the elements of context

In attempting to make a theoretical account of context, we have noticed that a large number of researchers have studied context related to a specific area or a specific problem, isolated from other context studies (see (Morris and Gruneberg, 1994)) for an overview of the experiments conducted for investigating context effects, in cognitive psychology). An important question is whether there may exist some aspects that are shared by several researchers or research communities that study context effects. A distinction between what may be referred to as *roles* and *elements* of a context, may help to answer this question. Studies show that the identification of context elements heavily depends on the type of task and domain in question. On the other hand, the role that context plays can, at an appropriate abstraction level, be generalized over specific tasks and domains. We will therefore start out by assuming that the role of context does not show much variance across domains or tasks, whereas the elements of a situation that play these roles do.

The essential aspects of context roles are captured by the notions of *relevance* and *focus*. Relevance refers to the usefulness of a solution to a problem in a particular environment. In problem solving, there exist essentially several lines along which to reason, and often, several alternative solutions to a problem. Context plays an important role in choosing the most relevant candidate. For example, recommending an angioplasty in a hospital which lacks the necessary instruments will not be useful. At a more detailed level, context is important for the generation and evaluation of explanations. When people ask why an airplane crashed, they will not be satisfied with the answer, 'because of gravity', even though this is not wrong. A possible acceptable answer should convey an anomaly that occurred, for example, in the engine of the plane. General, principled knowledge (text book knowledge) does not change across its users or the situations in which it is used. However, the *use* of principled knowledge is relative to the context in which it is applied. Relevance is, therefore, directly related to the *quality* of the solution produced for a problem.

The focusing aspect of context roles relates to the *efficiency* of the problem solving process *while maintaining relevance*. Generally, reasoner's face a huge search space. Efficient pruning of some parts from the search happens through the focusing of attention on particular regions of the memory, and as early as possible. Context serves as a focusing mechanism through determination of goals, and epistemological and physical needs of the reasoner, in order to accomplish the tasks that already active goals evoke.

Context elements are related to these two aspects of context roles, since the elements are the identified parts of knowledge and/or information that used by the reasoner to accomplish its task. At any time the attention of the reasoner is focused on a limited region of the domain knowledge. Contextual elements consist of a portion of that region. Therefore, in order to be able to refer explicitly to where the attention is focused, the contextual knowledge involved also needs to be referable. This requirement leads to a need for a context ontology.

3.2 A model of context - a high-level ontology

One of the first attempts to distinguish between context types is due to Hewitt (referred to in Tiberghien,1986). According to his distinction, context elements are classified into *intrinsic* and *extrinsic* ones. The intrinsic context elements are connected to the target item, while extrinsic ones are not actually part of the item in the memory. For example, the color of the wall of the room where subjects learned verbal targets is an extrinsic context element, while the color of the printed words read is an intrinsic context element. Hewitt's extrinsic context elements are predicted to assist recall, but not recognition. However, the findings of later experiments on verbal learning and face recognition contradicted these results, since they also detected context effects on recognition. Thus, Hewitt's context taxonomy fails to account for some of the empirical findings. An augmented and finer grained context typing is required in order to create a context account that can explain the tangled experimental results. One of the most important deficiencies in Hewitt's partial-ontology is that he did not consider the role of the 'state of the mind' of the reasoner in performing cognitive tasks. This amounts to not being able to cope with the differences in the nature and demands of diverse cognitive tasks, as it is the mind which ultimately selects the cognitive tasks to be fulfilled at each time.

Our starting point for a context ontology is that problem solving is a deliberate process in which two basic elements are the *agent* and the *external situation* in which the problem solving occurs. The key criteria for this distinction is the *deliberate* activity of a subject's mind. We do not study problem solving separate from the problem solver, as has traditionally been done in AI. In fact, the behaviour of an agent is shaped on the basis of two important factors: its own personal characteristic or state of mind, and the characteristics of the problem with which the agent is occupied. In order to reflect this aspect, we propose a distinction between *internal* and *external* context types, where internal context represents the 'state of the mind' while the, external one is related to the ground facts that happen to exist in a situation.

In our account, the agent's decisions regarding encoding details are shaped by his goals, hypotheses, and expectations, i.e. the internal context. Thus, the notion of internal context is a a step toward realizing the implications of 'level of processing' hypothesis, as well as of 'Encoding specificity' theory (see section 2).

External context, in turn, has two distinct groups of elements: those related to the *target* and those related to the *environment*. External context elements basically stay static during problem solving. That is, external context comprises the static facts in the problem solving situation. For example, in a clinical setting, the agent is the clinician (the reasoner), the target is the patient case, and the environment is the place where diagnosis and treatment occur. At the next level of specialization, the internal and external contexts are divided into *interactive* and *independent* types - in agreement with Baddeley's distinction. We also differentiate internal-context into interactive and independent, while Baddeley concentrates on external context. That is, he identifies what external context elements are independently or interactively encoded. The top-level ontology is shown in Figure 1. Here the two types of external context, target-related and environment-related, are modeled explicitly.

Notice the difference between our criteria and Hewitt's regarding the classification of context elements. In our work, the distinction between internal and external context is relative to the agent, while Hewitt's distinction is relative to the target. In attempting a mapping, our target-related context resembles Hewitt's intrinsic context while our environment-related one resembles his extrinsic context.

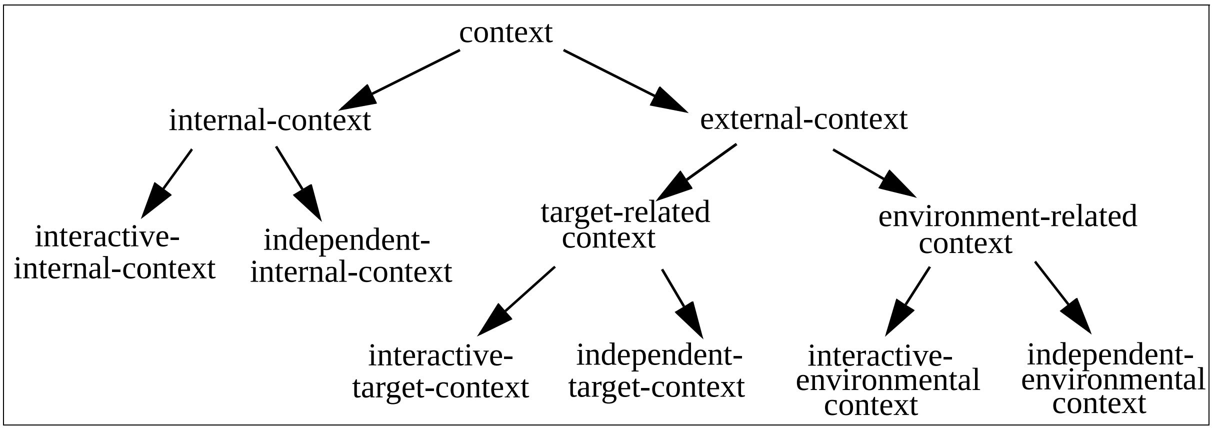

FIGURE 1. Context Ontology emphasizing the active role of the reasoner

Some examples may help to clarify the notion of context typing. First, the *independent* type: In medical settings, what the patient was wearing is of type target-related under external context, whereas the color of the examination room is of type environment-related. On the other hand, whether the clinician has short hair is of type "independent internal context". Regarding *interactive* context, for diagnosis and therapy tasks, goals and predictions (hypotheses) of the clinician are of type "internal-context", while the pregnancy condition or previous diseases of the patient are target-related. The characteristics (i.e., the conditions and constraints) of the place where the patient-clinician encounter occurred are of type environmental-related contextual aspects. For example, it is important whether the place is a well-established hospital or an emergency tent in a forest.

The ontology of top-level context elements just described is one step in establishing a model of context. The knowledge types should be related to the reasoning tasks to which they apply, and a suitable analytic method to study the relationships between knowledge types, tasks, and methods should be identified. This is precisely what a knowledge-level approach should provide.

3.3 A Task-centered approach to context focusing

Modeling the process of diagnosis involves identifying the subtasks which, when accomplished, lead to the achievement of the diagnostic goal, i.e., finding the fault causing the abnormal observations. It is exactly these subtasks that have a central role in modeling a shift of attention. This is because the subtasks are the *loci* where a shift in the focus of attention may be invoked. Knowledge level analysis opens up for starting out from different knowledge perspectives, e.g. a task perspective, a method perspective or a domain knowledge perspective. The tasks, in turn, identifies the need for knowledge, including contextual knowledge. We start by locating the subtasks of the whole diagnostic process where contextual knowledge is useful. After locating the points of contextual knowledge use, we try to understand the way these influences come into existence. The last step is in attempting to analyze the types of knowledge being utilized at these points.

Focus of attention is switched when the goal of a reasoner changes. A change in the goal signals that a new task is to be performed. Attention parallels the epistemological needs of the reasoner when engaged in a dedicated task. Therefore, a task is the natural medium for discussing shifts in the focus of attention. This is what we have adopted, leaning on the model proposed by van Heijst, Schreiber, and Wielinga (1997).

As earlier mentioned, two components pertinent to the context phenomenon are the state-of-mind and situational contexts. A very important issue is to identify and explicate the link between these two components of context notion. In our account, the internal context (corresponding to the "state-of-mind"), imposes the focus of attention which, in turn, is captured in terms of core domain concepts, cases, and the external context (see figure 2).

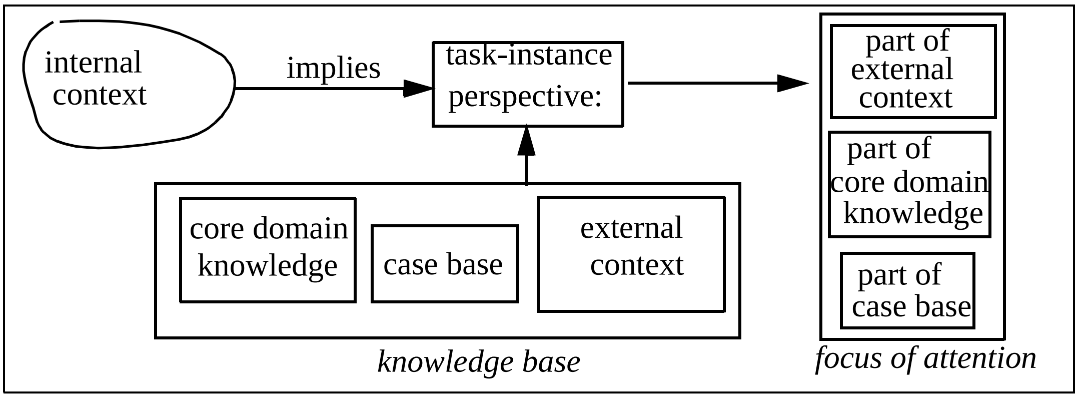

FIGURE 2. The relationship between two components of context

The modeling of contextual knowledge can be considered as consisting of two activities: (i) typing/categorizing contextual knowledge, and (ii) identifying what kind of role each type plays, and where/when this happens during problem solving, Implicitly suggested in these last two activities are the two perspectives that guide us toward an adequate account of the notion of context with respect to construction of knowledge-based systems. These perspectives are not mutually exclusive, but closely interrelated. They consider context, as respectively,

- a knowledge type
- a means for shifting attention of focus.

The first perspective leads to the study of how contextual knowledge can be categorized. So far context has usually been referred to as a unit, like a closed box. What we need is to look into this box, and break it into meaningful and useful pieces, in the form of knowledge types. The ultimate aim is to construct a context ontology and a contextual domain knowledge related to the application task of medical diagnosis. The role of context as a focusing mechanism is fundamentally important for complex tasks (Brezillon, 1993). The second perspective takes, as a starting point, an application task (in our case diagnosis) and identifies the subtasks of the overall diagnostic process where some contextual influences are anticipated, as well the role of contextual knowledge in each subtask.

The “attention of focus” perspective leads to the study of the relationship between a task and the type of the context it relies on. We may also say that the first perspective adopts a static view, while the second one is more concerned with dynamic aspects of context, i.e., its use.

For example, according to van Heijs et. al., (1997) the following activities lead to the construction of a knowledge-based system: (i) construct a task model for the application; (ii) select and configure appropriate ontologies, and if necessary refine these; (iii) map the application ontology onto the knowledge roles in the task model; (iv) instantiate the application ontology with domain knowledge. Although we have a different motivation (than reusability) for our study, we have found this account a useful basi for a knowledge level analysis of context. Our strategy can be expressed in terms of the following research subtasks:

- construct a model of diagnostic ‘process’ knowledge. This serve to explicate all the involved application subtasks (i.e., diagnostic subprocesses),
- identify the **loci** of contextual effects (i.e., the subtasks where context may have an influence)
- identify the **role** context plays in each locus (i.e., the way it effects each subtask)
- identify the **source** of context that can play these roles

3.4 The mechanism of shifting the focus of attention

Researchers from various communities studying areas such as the generation of explanations, user modeling in communication, natural language processing, and human-computer interaction have all acknowledged the importance of perspective (Suthers, 1991; Mittal and Paris, 1993; Pichert and Anderson, 1977; Lester and Porter, 1991). For example, regarding text comprehension, Pichert and Anderson argue that “if, for whatever reason, people take divergent perspectives on a text, the relative significance of text elements will change”. A presumption underlying the notion of perspective is that concepts are represented as a set of features, and not all the features of an item become activated each time the item is encoded. Each episode that contributes to the encoding of a concept refers only to a subset of these features. The task context in which a concept is presented determines that particular subset of encoded features. The importance of taking into account the type of task to be accomplished has recently been recognized (McCoy, 1989; Edmondson and Meech, 1993; Cahour and Karsenty, 1993).

Our account of the relation between the notions of context and perspective establishes a chain that starts at the goal. It is commonly agreed that goals are of crucial importance for contextualized problem solving and learning (e.g. Ram and Hunter, 1991; Leake and Ram, 1993; Bogdan, 1994). The goal determines the task necessary for its achievement. Once selected, the task then contributes to imposition of the perspective (see figure 3).

Task demands determine which aspects of a concept will currently be interactively processed, including both non-contextual and contextual features and cues. A coherent set of aspects constitutes a perspective. For example, the feature *weight* of the piano is not activated with respect to a going-to-piano-concert event, but it is activated in other contexts, such as when the piano is to be moved. So, each type of task needs only a portion of the domain knowledge for its accomplishment.

In our system, the task perspective (see figure 3) helps to determine the to-be-focused regions both of specific (episodic) and general domain knowledge. Related to the episodic part, it specifies the *types* of the episodes which are useful. For example, depending on the goals of the reasoner, either explanation-cases, or plan-cases are focused on, in our system. Task perspective specifies also which information is to be used as cues for episode activation. That is, it selects from the huge amount of available cues (both core and contexual cues) the ones which are currently relevant. Focusing on general domain knowledge happens through spreading activation. Task perspective determines the *upper bound*, the *lower bound* and the relevant *relations* for spreading activation.

As a result, internal context has a role in determining which features, including both external context and focal cues, that are relevant in a particular situation, and how relevant they are. In

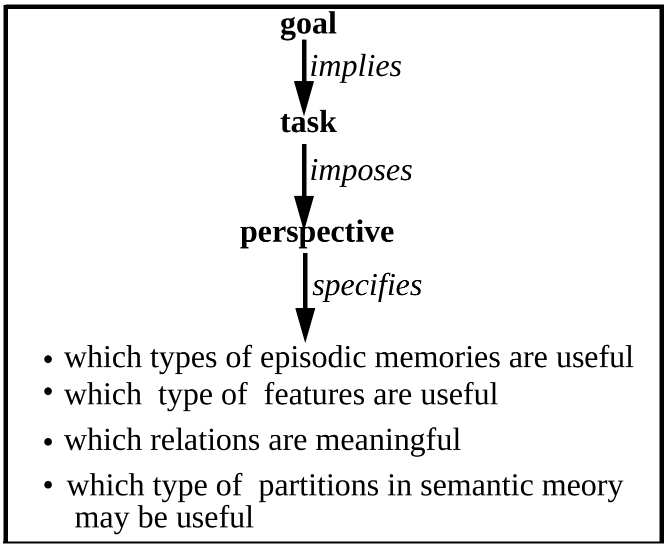

FIGURE 3. The explanation of how internal context serves as a focusing mechanism.

other words, internal interactive context determines the interactivity of external-context related features as well as non-context ones. Regarding external context elements, this influence is particularly important for tasks in which available information is incomplete, that is, focal retrieval cues are not sufficient to access target concepts unambiguously. In such cases, the context cues may strongly influence the degree of match between the present material and a past encoded material. Hence, internal interactive context is the main context type that imposes a perspective, which in turn determines the characteristics of both encoding and retrieval.

3.5 Abductive reasoning and experience

The expert’s problem solving in abductive domains, (for example, an expert physician’s diagnosis) has been argued to apply a ‘pattern recognition’ method (Patel, 1984) where the physician remembers past episodes of similar problems and use these episodes as short-cuts when solving new cases.

As Hatano and Inagaki (1992) assert, “the construction of models by abduction is nearly impossible when one does not possess a usable old model. People tend to be at a loss when they cannot think of any model that can explain the observed unexpected events”. According to them, it is hardly possible that we can formulate hypotheses which explain the surprising facts “unless we possess a familiar prototype” that we can take as the starting point. The notion of ‘familiarity’ (or family resemblance) is central to the later work of Wittgenstein (Wittgenstain, 1953), as well as the work of Pierce (1960). Pierce relates this notion to ‘abductive talent’, sometimes also referred to as instinct.

Educational psychology also supports this argument by proposing that memory experts is characterized by knowledge structures such as scripts and individual cases (Schmidt &Boshuizen 93). Their theory of “developing expertise” is based on the idea that the physicians become expert by learning from solving actual patient cases which they remember and utilize in later problem solving processes.

This view emphasizes the close interdependence between learning and remembering, which has also been stressed in cognitive psychology, in relation to the role of context in learning. An implication of this for AI is that a knowledge-based system exhibiting expert behaviour should be able to utilize past experiences. Case-based reasoning is nicely tailored for reflecting such behaviour. As context has been shown to have major influences on remembering episodes (Tulving 73), its inclusion in case-based problem solving is an improvement of the power of case-based approach.

Inspired by suggestions that an expert's hypothesis generation invokes a pattern recognition process (Patel, 1984), that medical experts (and others as well) construct memories from individual experiences (Schmidt &Boshuizen, 1993), and that what is retrieved by a recall activity is not only a stimulus but a particular experience of it (Tulving, 1973) we have adopted the view that abductive hypothesis generation involves a so-called 'cued recall'. Given a partial set of cues, the entire phenomenon is remembered. A good recall is conditioned on the similarity between the cues and the stored memories. Regarding hypothesis generation in particular, this idea may be conveyed as follows

> *Experiences, and ability to find similarities between the current problem and the previously experienced ones are the basis for the generation of good hypotheses.*

As Polya summarizes, "Analogy pervades all our thinking, our everyday speech and our trivial conclusions as well as artistic ways of expression and the highest scientific achievements" Polya, 1957).

Elstein (1995) states in connection with diagnostic problem solving in medicine that "'Intuition' includes the capacity to identify a puzzling case as an unusual presentation of a more familiar class". This indicates the importance of similarity notion in medical diagnosis. Apparently, the ability to recognize the similarity which may not always be detectable at the surface, but which requires deeper insight. This issue relates to Tulving's theory on the major effects of context on retrieving episodes from memory. The strong dependence between context and memory-retrieval success arise most probably from the role of context in similarity assessment of two episodes.

4 Context in medical diagnosis

On the ground of the results of the debate going on in the clinical reasoning research, and instructional research (studying the curriculum and different teaching methods), we view diagnostic problem solving as application of a hypothesize-and-test strategy. However this can be realized in different ways, and indeed is realized in different ways by different reasoners. For example experts are declared to use a 'pattern recognition method' (Patel 1984) which relies on past problem solving episodes, thus episodic memory (Schmidt and Boshuizen 1993). However, as typicality of a problem decreases, even high level experts use deep knowledge in increasing degrees. In this work, we assume an expert behaviour which primarily relies on a 'pattern recognition' method which we modeled as a case based reasoning method.

Two subtasks of diagnosis in which a clinician extensively uses context elements are generating hypotheses about the cause of a failure, and planning information gathering actions. Hypothesis generation is often achieved by retrieving past episodes where similar features have been experienced (Tulving et. al., 1973; Schank and Abelson, 1977). This task is related to the context dependency of memory, which we have elaborated in earlier sections. The other task for which a clinician utilizes context elements is the controlling and planning of actions. Observed from a distance, the clinician asks the patient questions, makes measurements, orders tests, prescribes medicine, etc. These actions are not arbitrary but controlled by external constraints. For example, in the case of deciding which of various therapeutic actions to prefer, the knowledge concerning which medicines have fewer side effects will play a role. Controlling and planning actions utilizes other types of context elements than when generating hypotheses. For example, whether the patient has been operated for a fault in the heart septum would probably not affect the selection of antibiotic to be used, but it is a rather relevant finding when generating diagnostic hypotheses.

In our clinical diagnostic problem solving model, two types of cases are considered: referred to as explanation cases and plan cases (see figure 4). The first is the encoding of the explanation that connects the relevant features to the disease. The second is the encoding of a plan that includes a sequence of diagnostic and therapeutic actions, such as defining an examination protocol, selecting the way to perform a particular test, recommending a treatment, etc.

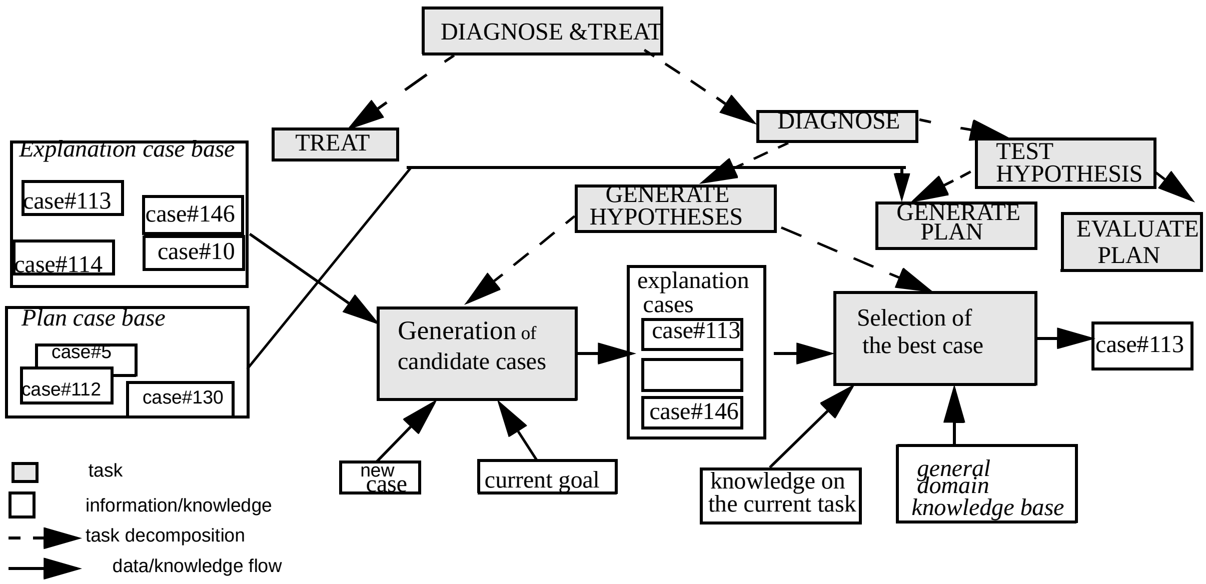

FIGURE 4. Main tasks of diagnostic problem solving and the use of experience in these tasks.

A part of the diagnostic task structure is illustrated in figure 4, related to an example run of the prototype implementation presented in section 5.1. Figure 4 shows the task hierarchy linked to some example cases used in that example.

To achieve a diagnostic goal, two subtasks must be invoked: Generating a diagnostic hypothesis, and testing (evaluating) the hypothesis. As shown in figure 4 cases of type explanation-case are used for hypothesis generation. They contain solutions that are regarded as diagnostic hypotheses. Plan cases are used in the hypothesis testing stage.

In diagnostic tasks, the target-related external context elements comprise *enabling conditions* for the development of a fault. Once established, the fault, in turn, leads to a set of consequences. In clinical reasoning, the patient's personal information constitutes the enabling conditions, while symptoms and signs reflect the consequences. Thus, the distinction between context and non-context features in a diagnosis task is based on whether the features existed independent of and before the development of the disease, or their existence should be explained by the disease.

Figure 6 shows the relationship between external-context, disease, and the consequences of the disease. The relationship between a disease and a focal-cue (i.e., consequence) is a 'causes' relation; that is, a disease CAUSES a consequence. This means that consequences become true after the development of the disease. On the other hand, the relationship between context elements and diseases are 'predisposes' and 'triggers'; a context-cue PREDISPOSES a disease or a context-cue TRIGGERS a disease. Predisposing and triggering context cues collectively constitute enabling factors. For example, context cues such as age, sex, and an earlier septum operation are predisposing factors for the disease infective endocarditis, while a recent dental operation can be a triggering factor. This is similar to the distinction between your leaving the door of your apartment open (predisposing factor), and a thief passing by your apartment (triggering factor).

Given our context model, we may enhance the domain model with an explicit representation of contextual elements. The relationship between three knowledge components' core-domain','external-context-domain´, and 'internal-context-domain' is shown in figure 5.

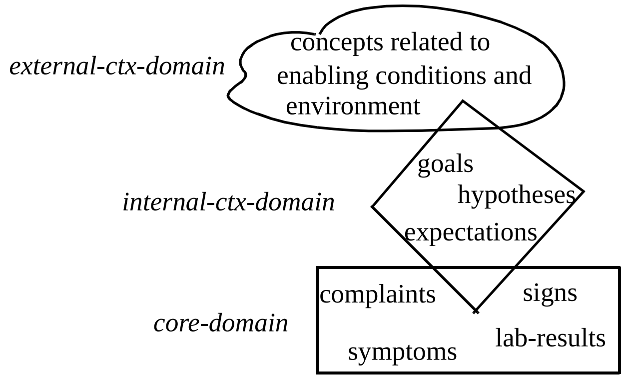

FIGURE 5. Domain knowledge, augmented with context elements

Figure 6 illustrates these three types of knowledge as different planes. *Core domain* knowledge consists of the knowledge needed to be able to come to *any* solution independent of time constraints and situation. This is the type of knowledge which has traditionally been studied. The *external-context-domain* contains knowledge related to external-context elements and usually consists of relevant aspects from several domains, while the *internal-context-domain* consists of elements of internal context which operationalize the core-domain as well as external-domain knowledge.

For example, in the medical context, biomedical knowledge is the core domain, while geographical and sociological aspects relevant to clinical reasoning would comprise external- context domains.These two planes are linked by *internal-context* elements, i.e., goals, and current and rejected hypotheses of the clinician.

An expert clinician very often uses her past experiences when encountering a new case. When finished with the case, she encodes the new case and retains significant parts of it in memory. In this reasoning model, the use of a past experience happens primarily through the retrieval of similar cases from memory. At the start, the available non-context cues are not sufficient for the clinician to remember a past case. Therefore, clinicians use various context cues available in the situation.

As shown in figure 1, external context is grouped into target-related (i.e. patient-personnel information) and environment-related information. In medical setting, patient-personal-information includes age, sex, occupation, economical conditions, earlier diseases, earlier operations, recent operations, family medical history, patient's emotions, concerns, level of intelligence, level of consciousness, ability to communicate, allergies, and habits. The majority of these act as enabling factors, as has been shown in figure 6.

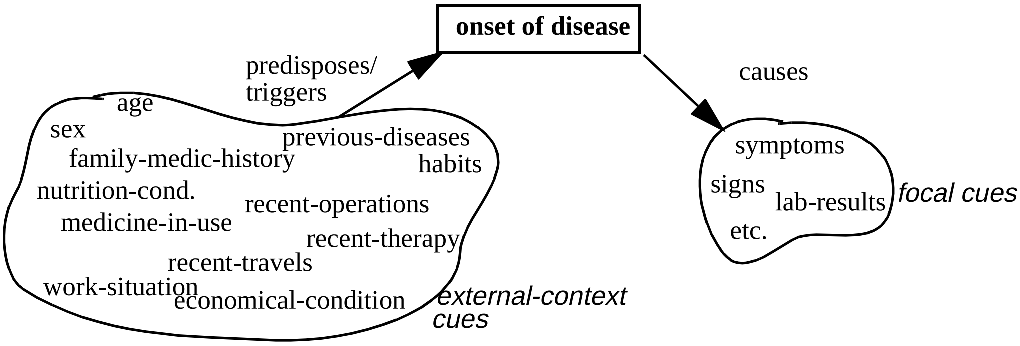

FIGURE 6. The boundary between target-related-context and non-context cues in medical domain

On the other hand, environmental information is related to equipment/tool availability, utilities/policies, time constraints, and other external constraints (e.g., spare beds). All elements of the external-context are the *ground facts* present in the situation. These are assumed to be static. Internal context, on the contrary, is more dynamic, and is usually changed by the reasoner during problem solving. It includes the goals of the reasoner, current predictions (hypotheses), failed predictions (e.g., ruled out hypotheses), and information needs of the reasoner.

4.1 Loci of context effects in diagnostic process

As we noted before, while internal context elements do usually display inter-process influences, the influences by external context elements can be investigated in the frame of separate processes, that is, local to processes. Our intention is to study external context in that order: first, locate the subprocesses where context effects can be investigated in the levels from which we can refer to our context ontology; second, identify which type of context element from that ontology are utilized in each subprocess, and finally analyze the nature of the interaction between the context elements and the subprocesses they are utilized in.

Figure 7 summarizes a portion of the subtasks underlying diagnosis which provides loci for possible context effects. The figure illustrates six distinct main loci for context assistance corresponding to six different stages of the diagnosis process.

subtask-loci	**role of context-** role	**type of context playing that role -** source
hypotheses activation	provides cues needed for activating most similar episodes	*external context*: primarily target-related ones, but also environmental context. These cues serve mainly as enabling factors in the disease process. *Internal context*: determines which types of hypotheses are to be activated as there may be various hypotheses (e.g., explanation-cases, or plan-cases).
hypotheses elaboration	pinpoints which features, aspects should be elaborated for deep similarity assessment	*internal context*: The task perspective determines the lower-bound (which includes elements from the *external context* (e.g., habits)), upper bound (which consists of the hypothesis for hypothesis elaboration), and the relevant relations for the spreading activation. The task perspective partitions also the cues (expressed as lower bound) into meaningful combinations illustrating which cues can be matched in order to make a similarity assesment contrained to meaningful comparisons.
strategy selection	gives possibility to compare alternatives	*internal context*: the set of *activated hypotheses* is studied to see whether there exists a hypothesis in the pool which will be wise to immediately try to either eliminate or confirm (e.g., life threatening). A strategy goal is determined.
hypotheses selection	provides criteria for deciding what to test	hypotheses selection happens according to the strategy goal chosen (*internal context*). For example, confirm the generated hypothesis A, eliminate hypothesis B, or discriminate between hypotheses A nad B.
prediction of consequences	modifies consequences (outcome) of a hypothesis	external context: mainly target-related. Modifies the outcome of the fault(e.g., disease) hypothesis (*internal context*). That type shows a synergistic effect on the consequence of a fault.
expectation testing	puts constraints on the actions. Determines which actions are possible, or wise.	*external context*: the resources in the environment (environmental context, e.g., available equipment), the condition of the target (e.g., patient-related context). For example special conditions of the patient that should be considered when selecting a test, or therapy.

FIGURE 7. Loci, role and source of context effects in diagnostic task.

5 Context in explanation-driven case-based reasoning

Our research group is pursuing methods for knowledge-intensive, explanation-driven case-based reasoning within open and weak-theory domains. This includes complex diagnosis and interpretation tasks (e.g. Grimnes and Aamodt, 1996). The context research reported here is done within the scope of (and to a large extent motivated by) this approach. The knowledge model in the underlying Creek system (Aamodt, 1994a) is a dense semantic network, where each node (concept) and each link (relation) in the network is explicitly defined in its own frame. A concept may be a general concept, a case, or a heuristic rule. It may describe domain objects as well as problem solving methods and strategies. The case-based method of Creek relies heavily on an extensive body of general domain knowledge in its problem understanding, similarity assessment, case adaptation, and learning. Cases, as well as general domain knowledge and information, are captured in a frame-based representation language (Aamodt, 1994b).

The underlying case-based interpreter in Creek contains a three-step abductive process of 1) activating relevant parts of the semantic network, 2) explaining derived consequences and new information within the activated knowledge structure, and 3) focusing towards a conclusion that conforms with the task goal. This “activate-explain-focus” cycle is a general mechanism that has been specialized for each of the four major reasoning tasks of the CBR cycle, as illustrated in Figure 4. The extensive, explanation-driven manner of utilizing general domain knowledge in the CBR suggests distinguishes Creek from most other CBR systems.

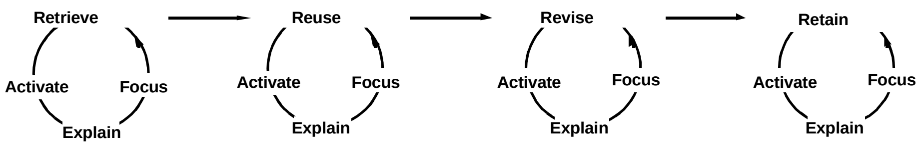

FIGURE 8. The CBR process and explanation engine

A Creek system has the potential to learn from every problem solving experience, either by storing the problem just solved as a new case, or by updating indexes. A new case is also created after a problem has been solved from rules or from the deeper knowledge model alone. The user is assumed to actively take part in both the problem solving and learning processes, e.g. by assessing hypotheses that the system cannot confirm or reject itself, supplying missing information, etc.

In our clinical diagnostic problem solving model, two types of cases are considered; explanation cases and plan cases. The first is the encoding of the explanation that connects the relevant features to the disease. The second is the encoding of a plan that includes a sequence of diagnosis and therapy actions, such as defining an examination protocol, selecting the way to perform a particular test, recommending a treatment, etc. A part of the diagnostic task structure is illustrated in figure 4.

5.1 Context in hypotheses generation

The inputs to the task of generating candidate cases are the information available in the new situation, including focal and contextual features, the case base, and the general domain knowledge base consisting of core and contextual domain knowledge bases. The output is a case providing a diagnosic hypothesis.

In the Creek framework, the diagnostician commutes between case-based and explanation-based reasoning methods, which utilize respectively, the episodic (case) memory and semantic memory, in order to generate hypotheses. The hypothesis-generation activity of the reasoner can be thought of as a “cued recall” in cognitive psychology terminology. One of the most prominent recall models in cognitive psychology is a two-staged model, also called “generate-and-test” model (Kintsh, 1970; Anderson and Bower, 1972). The original two-staged model takes into account only the deep knowledge, captured in a semantic memory. Recent models [Baddeley, 90) propose that the first stage of recall takes place in the episodic memory. In the Creek framework we model the cued recall (retrieval) as a three staged process, corresponding to the activate-explain-focus cycle, which has close similarities to the recent generate-recognize models. A difference is that, in our problem domain, we formulate a number of hypotheses suggested by past episodes, but focus only a smaller set of these. The reason is that, the past episodes to be retrieved are rather seldomly the identical copy of the initially provided case, contrary to the verbal learning or face recognition experiments where the identical episodes are remembered. We do not retrieve identical episodes, but the more or the less similar ones. Conseqeuntly, we have an additional stage where we select between these retrieved, similar cases. The threestages of our recall corresponds to the CBR tasks of activate, explain, and focus. The recall of past episodes is realized as a explanation-based reasoning supported case-based reasoning process.

We conceive hypothesis generation (i.e., cued recall) as including the phases of activating plausible hypotheses, elaborating them, as well as of selecting the "best" one for closer scrutiny, that is, for testing. This is an abductive process as the hypotheses so generated may be rendered incorrect in the testing phase.

In diagnostic problem solving, activation involves remembering previously experienced *similar* problems. Since similarity assessment, in this stage, is based on associational resemblance, the similarity match happens rather quickly but may give superficial results. In fact, a stored episode may either seem different at first sight but be rather similar at the deeper level, or it may seem more similar at the surface than it indeed is at dept. Therefore, the activated cases enter into an elaboration phase where their similarity with the new case are more thoroughly assessed. So, the elaboration process (explain, in Creek terminology) utilizes the deep knowledge in the semantic memory.

The very first question that comes to mind is whether context information is needed when activating (i.e., formulating) hypotheses in the first place, or whether it only applies when elaborating the activated hypotheses and selecting the most promising one. This is analoguous to debating whether all the meanings of an ambiguous word are activated on encountering the word, or the activation is constrained by context.

In our model, context assists both. It is important to notice, however, different types of contextual knowledge and information are effective in the two stages. The influence of context in activation is important for the efficiency of abductive inference since its efficiency is proportional to the number of hypotheses activated. Contextual information should have a constraining effect on the activation of hypotheses. The smaller the number of activated hypotheses, the more effective the abductive inference will be, under the condition that the activated set includes the best hypotheses. So, context should not constraint the activation for any price; the activated set should embrace the most relevant hypotheses as well.

The heart of this account is that both independent and interactive type of contexts can increase the distinctiveness of an episode. Distinctiveness corresponds to being easily differentiable from alternatives. This implies that independent context is also effective in the activation stage, as this stage relies on shallow associations, and does not necessarily require interactive encodings. Weak domains such as medicine are usually less well-understood. Sometimes, the physicians feel there is a relation between two entities but they are not able to explain the relationship. For example, they may have statistically noticed that in a particular illness, the patients are very often male. They cannot embody this knowledge in deep knowledge, but this may be stored as independent context, and thus used in activation stage.

On the other hand, the elaborative similarity association needs deep knowledge, and thus interactive encodings. Consequently, only interactive context is effective in the elaboration phase. We also view elaboration as a justification of similarity. This is inherent in abductive reasoning, since it is characterized by fallibility, its reasoning needs be justified, even though this is different from verification of the ultimate solution/result. This model is an example of integrated reasoning behaviour where one reasoning paradigm is supported by another. In our case, case-based reasoning is supported by explanation-based reasoning when the reasoner needs to fall back on deep knowledge. The elaboration allows ranking the retrieved cases with respect to their ‘explained’ similarity to the new case. This is a simple step in our model, and will not be further discussed here.

The last stage of our recall model is focus, where the case which is most similar to the new case is chosen among the recalled ones.

Hypothesis activation:

A past case is plausible to activate if its solution can be utilized as a ‘hypothesis’ for the new case. The hypothesis should explain the findings. However, as the domain is characterized by ill-described problems, not all relevant findings are known at the outset. Often, there is other information - contextual information - avaible. For example, in the generation of diagnotic hypotheses, what we call ‘enabling conditions’ predispose or trigger the onset of the disease. So, its role in a disease process comes before the establishment of the disease. As such, it has no impact on the explanation chain between a disease and its consequences, the findings. This contextual information is not to be explicitly explained, but contributes to the explanation of to-be-explained findings indirectly, by explaining the existence of the disease. Therefore, starting from contextual information as well as from findings, we may infer the disease more certainly. Figure 9 illustrates how the activation of hypotheses from the case base is conceptualized.

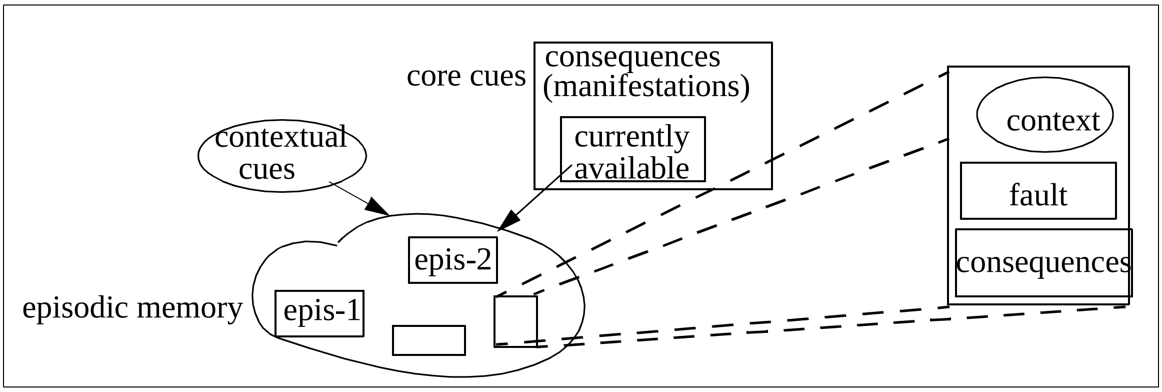

FIGURE 9. Activation of hypotheses

The type of contextual information which is relevant for the activation process is described in the slot ‘relevant-case-to-entity-relations’ in the dedicated task frame This is an the *external* type of context. Examples are patient-related contexts which works as enabling factors for a disease and specified via the slot names such as has-age, has-recent-disease, has-recent-surgery, has-sex, has-occupation, has-econ-sit, has-psyc-sit, has-prev-disease, has-habit, has-family-hist, has-recent-therapy, etc. Which cases that may be activated is based on the features determined by the ‘relevant-case-to-entity-relations’ and the available findings.

The set of relevant features determines the combined reminding. The past cases of which combined remindings reach a predefined threshold are activized. This activation process occurs in the episodic case base. Both interactive and independent contextual features may be relevant for activation. So, the entire situation is utilized as the basis for the activation process. In this way a number of cases providing plausible hypotheses are activated.

Figure 10 is an excerpt from an example run of the system developed to experiment with different context models. Three type of cases are shown (a new problem case and the two types of stored cases). The output sequence illustrates hypothesis generation and selection, as described in below.

Hypothesis elaboration - explanation: The selection of the most plausible case is based on the similarity judgment of candidate cases. Similarity of two cases can be exposed directly or indirectly. Direct match occurs when two cases have identical slot-value pairs, e.g, both cases having has-sex:female. Hypothesis activation in the 'activation' phase relies on such direct matches. These embody superficial similarity. In 'explain' phase we try to catch deeper similarities which are not visible at first sight.

An indirect match happens when a path between slot values of the two cases is found in the general domain knowledge. In the example (see figure 10), the information has-occupation:unemployed in the new case has-habit: drug-abuse in *case#113* are found indirectly to weakly match, as explained by the path 'drug-abuse *causes* immunosupression *caused-by* malnutrition *caused-by* bad-econ-sit *associated-with* unemployed'. This is illustrated by the explanation structure at the bottom of the figure. This explanation is relevant in the context of the hypothesis endocarditis, established earlier in the example run. The existence of such paths contributes to the similarity of two cases. The strength of this similarity depends on the strength of the relations connecting the concepts on the path, as well as the length of the path. So, the indirect matching process involves a search in the general domain knowledge that investigates whether apparently dissimilar features could be similar through a finer grained matching process. The task knowledge, together with the generated hypotheses, set up a context which helps to identify the portions of the general domain knowledge that are relevant for search. As shown in the beginning of the example run, the explanation cases *case#113, case#146, case#114* are candidate cases.

The process of selecting the best case involves thoroughly matching each of these cases with the new case. During the evaluation of *case#113*, the values of the slots has-recent-surgery, has-previous-disease, has-recent-therapy and has-habit are identified by the task to constitute a meaningful partition in the domain knowledge in order to carry out similarity match. Partitioning the fetures into meaningful sets can alleviate the the explosion of similarity matches. In order to constraint the search for similarity, we focus on similarities within partitions. For example, an attempt to match has-recent-surgery:dental-surgery in *case#113* and has-recent-therapy:acupuncture in the new case is allowed while matching has-age:22-years-old in *case#113* and has-recent-therapy:acupuncture in the new one is prohibited by the task knowledge. It is the nature of the task, that is, the perspective imposed by the task which determines what can be a meaningful partition. The slot 'has-meaningful-partition' lists the groups of features which constitute a partition.

The similarity elaboration takes place in the semantic memory. A focused spreading activation picks out the portion of that memory to work on. The intersection of partitions constitutes the lower-bound for spreading activation. The upper-bound is also specified in the task frame. It is 'hypotheses' for generating diagnostic hypotheses. The activated hypotheses instantiate this role. The other element of the perspective provided by the task frame is the set of meaningful relations along which the spreading activation takes place. The first explanation structure near the bottom of the figure shows that has-age can only match has-age relation while the other two structures illustrating explanations matching relations has-recent-surgery, has-recent-therapy, has-prev-disease and has-habit. That is, the task knowledge does not specify a partition having the slot combination has-age and has-recent-surgery as a legal one. The local context that defines the legal slots for matching are listed in the :slot part of the explanation structures.

```
NEW-CASE
has-sex              female
has-age              30-years-old
has-occupation       unemployed
has-prev-disease     aortic-valve-disease
has-recent-therapy   acupunctur
has-findings         low-bp  sweating
has-location         well-established-hospital
```

```
case#113
has-type: explanation-case
has-sex: female
has-age :22-year-old
has-prev-disease: calcific-aortic-stenosis
has-recent-surgery: dental-surgery
has-habit: drug-abuse
has-findings : low-bp   high-fever   fatique
               high-sedimentation  vegetation
has-solution: endocarditis
```

```
case#130
has-type: plan-case
gathers-info:  high-fever fatique
has-location :  well-established-hospital
control-schema:  action-21  action-36
 . . . . .       action-90   action55
```

EXCERPT FROM OUTPUT:

```
Starting Problem Solving. .
GOAL: g-diagnose

Starting the task Generate Hypotheses..
Starting the task Generation of Candidate Cases..
Candidate cases are:  case#146, case#113, case#114

Starting the task Selection of the best case..

Evaluation of case#113:

Hypothesis is Endocarditis
 . . . .
Attempting to find path from DENTAL-SURGERY to AORTIC-VALVE-DISEASE
Attempting to find path from DENTAL-SURGERY to ACUPUNCTUR
Attempting to find path from CALCIFIC-AORTIC-STENOSIS to ACUPUNCTUR
Attempting to find path from FARMER to UNEMPLOYED
Attempting to find path from 22-YEAR-OLD to 30-YEAR-OLD
Attempting to find path from DRUG-ABUSE to UNEMPLOYED
 . . . .
The result of Similarity Match between the new case and case#113:
((#S(Expl-struct
  :SLOT  (HAS-AGE)
 :FROM 22-YEARS-OLD to 30-YEARS-OLD
 :METHOD. .
   (22-YEARS-OLD  INSTANCE-OF  ADULT HAS-INSTANCE 30-YEARS-OLD)

((#S(Expl-struct
 :SLOT  ( HAS-RECENT-SURGERY HAS-RECENT-THERAPY HAS-PREV-DISEASE HA
 :FROM DENTAL-SURGERY    :TO ACUPUNCTUR
 :METHOD
   (DENTAL-SURGERY          INSTANCE-OF
   BACTERIEMIAE-PROVOKING-FACTOR  HAS-INSTANCE ACUPUNCTUR)

((#S(Expl-struct
 :FROM DRUG-ABUSE  :TO UNEMPLOYED
:SLOT  ( HAS-RECENT-SURGERY HAS-RECENT-THERAPY HAS-PREV-DISEASE HAS
 :METHOD
  (DRUG-ABUSE  CAUSES  IMMUNOSUPRESSION  CAUSED-BY  MALNUTRITION  C
  BAD-ECON-SI T   ASSOCIATED-WITH   UNEMPLOYED )

   . . . . . . . . .
```

FIGURE 10. An excerpt from example run.

The hypothesis, that is the solution of case#113, in this example is endocarditis. In the general domain model (not shown) endocarditis is associated with being adult or elderly. The patients of both new-case and case#113 have been inferred to be adults during similarity matching. This indicates that these two cases are similar regarding age, in the context of endocarditis. This may not be so in contexts of other diseases. Therefore, when matching similarity of cases, a relevant region in general domain knowledge is identified as having concepts and relations circumscribed by a hypothesis and a partition which is specified by the task.

5.2 Hypothesis testing

The solution of the best case is now the hypothesis which will be evaluated. A typical characteristic of new cases in a medical setting is that they contain incomplete information. Therefore, an information-gathering process begins, in order to determine whether other consequences of the hypothesis (the disease) are also present in the new situation. The retrieved case guides the information-gathering process. The goal is now to collect data on the findings on which the retrieved case has information while the new one does not.

The task of evaluating hypotheses invokes two subtasks: generate plan task, and evaluate plan task. Generation of plans occurs by retrieving past plans from the case base. The process of plan evaluation involves executing the retrieved plan. A plan is indexed with the 'information goals', i.e, features to be gathered. These features refer to the consequences expected to be present if the hypothesis (the disease) is correct. For example, assume that the retrieved explanation case has solution endocarditis (current hypothesis) which is expected to cause low blood pressure, and that the new case lacks information about blood pressure. One of the information goals becomes ' blood pressure ', and a plan to gather this information, among others, will be retrieved. A plan case also includes the set of actions to take in order to gather such information. Another part of plan indices consists of external context elements that may influence the action selection. For example 'location' captures information about the available resources in the environment. When the location is a 'well-established hospital', in order to manage bleeding, two alternative actions, 'blood transfusion' or 'use packed red blood cells', are both possible. In a village, on the other hand, packed red blood cell may not be available. Thus, for a useful plan retrieval, plans are indexed by location, among other external context elements. An example plan case is shown as case#130.

Notice that in plan-case retrieval as well as in explanation-case retrieval, we utilize interactive external context elements only, and not independent ones, despite the findings from empirical studies in cognitive psychology that indicate on independent context assistance when generating candidate episodes (cases). The reason is that the important point in clinical problem solving is not to remember complete patient cases, including information of type independent patient context such as skin color, but to remember a patient showing similar signs and symptoms related to the same enabling conditions. Furthermore, this type of context (e.g., independent-target-context) information will generally be different in new and past cases. Consequently, taking these context elements into consideration would mean deliberately decreasing the retrieval performance. Knowing this and being able to decide which context elements should effect the retrieval (people may not have this possibility as this may partly be an unconscious decision), we prefer to ignore context elements that will negatively impact the retrieval performance. Therefore, independent-target-context is ignored.

6 Summary and Conclusion

In this paper we have described a theoretical account of context modelling, related to the abductiive diagnostic task. Our emphasis has been on establishing the necessary basis for our model, which has been based on earlier context studies in AI as well as cognitive psychology.

A knowledge-level approach was chosen for the analysis of interrelations between context as a knowledge type, reasoning tasks involving context use, and problem solving strategies.

As system complexity grows, and the demands for flexibility and adaptability steadily increases, the AI community needs to focus on methods that enable decision-support systems to provide the functionality wanted. We have shown a way in which an explicit model of context can be captured and made useful in a hybrid case-based and explanation-based reasoning system. The combination of case-specific knowledge with more general contextual knowlegde seems particularly promising in 'grounding' our systems better to the real world.

For adatptive systems, the learning issue is of course important. It has not been addressed specifically here, since contextual learning within our framework is part of ongoing research, as is also more extensive testing and evaluation of the method decsribed here.

7 References

Aamodt, A. (1991) A knowledge-intensive, integrated approach to problem solving and sustained learning, Ph.D (Dr.Ing.) dissertation, University of Trondheim, Norwegian Institute of Technology, Department of Computer Science, May 1991.University Microfilms PUB 92-08460, 1992.

Aamodt, A. (1994a). Explanation-driven case-based reasoning. In S. Wess, K. Althoff, M. Richter, Eds. Topics in Case-based reasoning, pp 274-288. Springer Verlag.

Aamodt, A. (1994b). A Knowledge Representation System for Integration of General and Case-Specific Knowledge. Proceedings from IEEE TAI-94, International Conference on Tools with Artificial Intelligence. New Orleans, November 5-12, pp 836-839.

Abu-Hakima, S. & Brezillon, P. (1995). Principles for the application of context in diagnostic problem solving.In Brezillon, P. and Abu-Hakima, Eds. Working notes of the IJCAI-95 workshop on "Modelling context in knowledge representation and reasoning".

Anderson, J. R. and Bower, G. H. (1973). Human associative memory. Winston.

Anderson, J. R. and Bower, G. H. (1972). Configurational properties in sentence memory. Journal of verbal learning and verbal behaviour. Vol 11.pp 594-605.

Baddeley, A. (1982a). Depth of processing, context, and face recognition. Canadian Journal of psychology . Vol. 36, No. 2, pp. 148-164.

Baddeley, A. (1982b). Domains of recollection. Psychological Review. Vol. 86, no.6, pp.708-729.

Baddeley, A. (1990). Human memory. Lawrence Erlbaum.

Barrows, H. S. (1994) Practice-based learning: problem-based learning applied to medical education. Springfield.

Bogdan, R. J. (1994). Grounds for Cognition: How goal-guided behaviour shapes the mind, Lawrence Erlbaum associates.

Breuker, J. and Van de Velde, W (1994). CommonKADS library for expertise modelling; reusable problem solving components. IOS Press.

Brezillon, P. (1993) , Proceedings of the IJCAI-93 Workshop on "Using knowledge in its Context", Laforia

Bylander, T. and Chandrasekaran, B. (1988) Generic tasks in knowledge-based reasoning: The right level of abstraction for knowledge acquisition. In Gaines, B. R. & Boose, J. H. Eds. Knowledge acquisition for knowledge based systems, vol. 1. pp 65-77. London: Academic press.

Cahour, B. and Karsenty, L. 1993. Context of dialogue: A cognitive point of view. In Brezillon, P. and Abu-Hakima, Eds. Working notes of the IJCAI-95 workshop on “Modelling context in knowledge representation and reasoning”, pp 20-29

Chandrasekaran, B. (1987). Towards a functional architecture for intelligence based on generic information processing tasks. In Proceedings of the 10th IJCAI. pp 1183-1192. Milan,Italy.

Chandrasekaran, B. (1992). Task-structure analysis for know-ledge modeling. Communications of the ACM, Vol. 35, no. 9, Sep. 1992, pp 124-137.

Craik, F. I. M. and Lockhart, R. S. 1972. Levels of Processing: A framework for memory research. In Journal of verbal learning and verbal behaviour. Vol.11, pp.671-684.

Craik, F. I. M. and Tulving, E. (1975). Depth of processing and the retension of words in episodic memory.Journal of experimental Psychology: general. Vol. 104-3, pp. 268-294.

Dojat, M. &Pachet, F. (1995). Three compatible mechanisms for representing medical context implicitly. In P. Brezillon & S. Abu-Hakima, Eds. Working notes of the IJCAI-95 workshop on “Modelling conetxt in knowledge representation and reasoning. pp. 179-182. Laforia.

Edmonson, W. H., and Meech, J. F. (1993), A model for Human-Computer Interaction. In Brezillon, P. and Abu-Hakima, Eds. Working notes of the IJCAI-95 workshop on “Modelling context in knowledge representation and reasoning”. pp. 31-38.

Eich, J. E.(1980), The cue-dependent nature of state-dependent retrieval. Memory & Cognition. Vol.8, pp. 157-173.

Elstein, A. S. (1995), Clinical reasoning in medicine. In J. Higgs, & M. Jones, Eds. Clinical reasoning in the health profession. pp 49-59, Oxford: Butterworth Heineman, .

Fisher, D. and Yoo, J. P. (1993). Categorization and problem solving, The psychology of learning and motivation, vol. 29, pp 219-255.

Godden, D.R. and Baddeley, A. D. (1975) Context dependent memory in two natural environments: On land and underwater. British Journal of psychology. Vol. 66, pp. 325-331.

Godden, D. and Baddeley, A. (1980), When does context influence recognition memory?British Journal of Psychology. Vol. 71, pp. 99-104.

Grimnes, M., Aamodt, A. (1996). A two layer case-based reasoning architecture for medical image understanding. In Smith, I., Faltings, B. Eds. Advances in case-based reasoning, Springer Verlag, Lecture Notes in Artificial Intelligence 1168. pp 164-178.

Hatano, G. and Inagaki, K. Desituating cognition through the constuction of conceptual knowledge. In P. Light and Butterworth, G., Eds. Context and Cognition. pp. 115-133. Harvester Wheatsheaf.

Haugeland, J. (1984), First among equals. In W. Kintsch, J. R., and P. G. Polson, Eds. Methods and Tactics in Cognitive Science. pp 85-99. Lawrence Erlbaum Associates.

Kintsch, W. (1974). The represenation of meaning in memory. Lawrence and Erlbaum.

Kintsch, W. (1970). Learning, memory and conceptual processes. New York: Wiley.

Kokinov, B. (1995). A dynamic approach to context modeling. In Brezillon, P. & Abu-Hakima, S. Eds. Working notes of the IJCAI-95 workshop on “Modelling conetxt in knowledge representation and reasoning. pp. 199-209. Laforia.

Leake, D. and Ram, A. (1993), Goal-driven learning: Fundamental Issues and Symposium report . AI Magazine, winter.

Lester, J. C., Porter, B. W. (1991), Generating Context-sensitive explanations in interactive knowledge-based systems, AI 91-60 , University of Texas.

McCoy, K. F. (1989). Generating context-sensitive responses to object-related misconceptions. AI Journal , vol 41, pp 157-195.

Mittal, V. O., and Paris, C. L. (1993). Context: Identifying its elements from the communication point of view. In P. Brezillon & S. Abu-Hakima, Eds. Working notes of the IJCAI-95 workshop on "Modelling context in knowledge representation and reasoning. pp 87-97. Laforia

Morris, P. and Gruneberg, M.(1994). Theoretical aspects of memory, London: Routledge.

Musen, M. A. Automated support for building and extended expert models. Machine Learning. Vol. 4. pp.347-376.

Newell, A. (1982). The knowledge level, Artificial Intelligence, Vol. 18.pp 87-127.

Patel, V. (1984). Cognitive processes in comprehension and knowledge acquisition by medical students and physicians. In SSchmidt, H. G. and de Volver, M. L., Eds. pp . 141-157. Assen: Van Gorcum.

Peirce, C. S. (1960) Collected papers. The Belknap press of Harvard university press.

Pichert, J. W., & Anderson, R. C. (1977). Taking different Perspectives on a story. Journal of Educational Psychology, Vol.69, No. 4, 309-315.

Polya, G.(1957). How to solve it: A new aspect of mathematical method. New York: Princeton university press..

Ram, A. and Hunter, L. (1991). A goal based approach to intelligent Information retrieval, in Machine Learning: Proceedings of the 8. International Conference, pp 265-269.

Schank, R. and Abelson, R. P. (1977). Scripts, plans, goals and understanding. An inquiry into human knowledge structures. NJ: Erlbaum, Hillsdale.

Schmidt, H. G. and Boshuizen, H. P. A. (1993). On acquiring expertise, Educational psychology review, vol. 5, No. 3, pp. 205-221.

Smith, S. M. (1986). Environmental context-dependent recognition memory using a short-term memory task for input. Memory & Cognition. Vol. 14, No. 4, pp. 347-354.

Steels, L. (1990). Components of expertise. AI Magazine. Vol. 11. No 2. pp.30-49.

Steels, L. (1993). The componential framework and its role in reusability, In J-M. David, J-P. Krivine, R. Simmons (eds.), Second generation expert systems. Spinger. pp 273-298.

Suthers, D. D. (1991), A task-apropriate hybrid architecture for explanation. Computational Intelligence. vol.7, pp 315-333.

Tiberghian, G. (1986) Context and cognition: Introduction. Cahiers de psychologie cognitive. European bulletin of cognitive psychology. Vol.6, No. 2, pp. 105-119.

Thomson, D. M. and Davies, G. M. (1988). Memory in context: Context in memory. Wiley.

Tulving, E. and Thompson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. Psychological Review. Vol. 80. pp. 353-370.

Turner, R. (1994). Adaptive reasoning for real-world problems, Lawrence Erlbaum associates.

van Heijst, G., Schreiber, A. Th., and Wielinga, B. J. (1997). Using explicit ontologies in KBS development. International Journal of Human-Computer Studies. Vol. 46, pp 183-292.

Watkins, M. J., and Tulving, E. (1975), Episodic memory: When recognition fails.Journal of experimental psychology: General. Vol. 104., No.1, 5-29.

Wickens, (1987) The dual meaning of context: Implications for research, theory , and applications. In D.S,. Gorfein and R.R. Hoffman, Eds. Memory and learning. Lawrence Erlbaum associates.

Wielinga, B.J., Schreiber, A.T., and Breuker, J.A. (1992). KADS: A modelling approach to knowledge engineering. Knowledge acquisition. vol.4, no.1. pp 5-53.

Wittgenstein, L. (1953), Philosophical investigations. Blackwell, 1953. pp 31-34.

Özturk, P. and Aamodt, A.(1997). Towards a model of context for case-based diagnostic problem solving. In Context-97; Proceedings of the interdisciplinary conference on modeling and using context. Rio de Janeiro, February 1997. pp 198-208.