Towards a Model of Context for Case-Based Diagnostic Problem Solving

Pinar Öztürk and Agnar Aamodt
Department of Informatics, Norwegian University of Science and Technology
N-7055 Trondheim-Dragvoll, Norway
Phone: +47 73598717 Fax: +47 73591733 E-mail: {pinar, agnar}@ifi.ntnu.no

Abstract: This paper presents a model of context based on the *roles* and *elements* of various context types. Two important roles of context are related to the notions of *relevance* and *focus*. The former is important for the *quality* of the results reached by a problem solving or learning task, while the latter is important for the performance *efficiency* of the task. Problem solving can be viewed as search in a large problem space where search for different entities is invoked at different stages. Context has a pruning effect on search, increasing proportionally to the incompleteness of the information at hand. Depending on the type of memory structure to be searched for, different types of contexts assist the access to memory. We attempt to give an account of the various types of contexts that facilitate memory access and utilization for different type of tasks. The criteria for distinguishing between several types of context elements are presented, and a context ontology based on these criteria is suggested. We then show how this account is integrated with a case-based approach to clinical problem solving.

1 Introduction

A strong motivation for research on context is its important role in information processing. It is generally agreed that context facilitates the selective processing of information, whereas disagreement exists as to what context actually is, and in what way it affects processes of problem solving and learning. An earlier trend was to consider context as a wholistic phenomenon, and to review the context effects on various processes in a uniform way. This entails referring to context without making a distinction between various types of contexts. Yet, different types of context elements have been observed to indicate significant difference in their 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.

The work described in this paper is an attempt to clarify the elements of context relevant to a problem situation, and the roles of these elements in constraining and guiding a reasoning process. Our general research agenda is to improve artificial intelligence methods for problem solving and learning in open and weak-theory domains. An example of such a domain is medical diagnosis and treatment. Since there exists no neat theory from which to deduce conclusions for these type of 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. For this type of domain and inferencing, 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 91]. 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. A term often used in relation to context is “perspective”. Perspective, in our account, is the set of relevant aspects one takes into consideration when accomplishing a particular task. The use of context to identify appropriate perspectives, and the use of perspective as a kind of filter or lens into a multi-task, multi-perspective knowledge base, is a core mechanism in our approach.

In the next section we give an account of how we assess the role of context at a general level, i.e. in a task-independent way. This is followed in section 3 by an identification of possible context types and elements, as suggested by cognitive science research. In section 4 we present a context ontology that makes explicit the different types of context elements relevant to problem solving and learning. Section 5 discusses perspectives, while in section 6 a model of context related to medical diagnosis and treatment is presented. In section 7 we describe how this model may be used to constrain and focus knowledge-intensive case-based problem solving and learning in this domain. The final section discusses the results and points out future research issues.

2 The role of context - relevance and focus

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. 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 the elements of the context and the roles of the context may help answering 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 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 role do.

The notions of *relevance* and *focus* capture the essential aspects of context roles. 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. People asking why an airplane crashed 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 proportional to the *quality* of the solution produced for a problem.

The other main role of context in problem solving and learning is its focusing ability. Focus is important, for ensuring *efficiency* of the problem solving process while maintaining relevance. At any point, the attention of a person is focused on particular issues and aspects.

Generally, people face a huge search space. Pruning of some parts from the search happens through focusing the attention only to particular regions of the memory, 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 active goals evoke.

In this work, we put a special emphasis on the influence of context on two important reasoning tasks: *memory use* and *action planning*. Memory is of crucial importance for intelligent behaviour. It was, in fact, memory studies that alerted some influential researchers to the importance of context [Tulving 73, Eich 80, Baddeley 82a-b, Thomson 88]. The retrieval process undergoing remembering may involve free-recall, cued recall or recognition, depending on the number of available cues and on how specific these cues are. Memory is said to be context-dependent when context improves the retrieval process. The task of action planning involves interaction with the real world. For example, in a medical setting, the diagnosis process includes information gathering actions such as measuring blood pressure and taking X-rays. The process of action planning should consider the resources available, as well as the policy to be adopted in a particular situation.

3 The elements of context - psychological studies

Despite the fact that cognitive psychology has mainly concentrated on studying contextual effects for tasks such as verbal learning and face recognition, these studies can assist our investigation of contextual effects in problem solving.

The initial studies of retrieval and the relationships between recall and recognition suggested context effects on retrieval, and context has been one of the primary mechanisms used in memory theories. [Anderson 73, Kintsch 74, Thomson 88]. The results of experiments, however, agreed on the existence of context effects in retrieval but differed in their account of how context influenced recall and recognition.

A number of verbal-learning researchers reported context effects in recall, but not in recognition. These studies arranged experiments where words were learned and tested either in the same or different contexts. For example, learning happened on land and was tested under water [Godden 75], or learning and remembering was tested in one room versus another (see [Smith 86] for review). The results announced that the reinstatement of the context resulted in better recall but did not influence recognition. Another type of experiment which has drawn considerable attention is state-dependent learning. In these experiments, the internal state of subjects are induced by alcohol or pharmacological means. Results show that state-dependent learning is observed to some degree on recall but not on recognition [see Eich 80].

One of the first attempts to distinguish between context types is due to Hewitt (referred to in Tiberghien [86]). According to his distinction, context elements can be classified into intrinsic or 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 another group of experiments on verbal learning contradicted these results, since they also detected context effects on recognition. For example, subjects were shown sentences containing a noun and adjective pair at learning time. At recognition time, they were asked to decide whether the nouns in the sentences were seen at learning time. Some of the nouns were paired with the same adjectives as at the learning time, some were paired with different adjectives. The nouns paired with the same adjectives at learning time was better recognized. An example is the sentence" The CHIP DIP tastes delicious". At retrieval time the recognition of CHIP in a sentence containing CHIP DIP (a kind of sauce) pair was superior to when the pair was SKINNY DIP (a way to swim).

Other empirical findings that show context effects on recognition come from face recognition and eyewitness studies. In a number of studies, the face learning process is modified by instructions such as 'assess the honesty of the face' or 'assess the relationship of that face with the one in the background'. Experiments showed that the changes in context such as face associations and instructive cues resulted in decreased recognition accuracy, at test time.

Theories have been developed to explain such a variety in the empirical findings. Tulving [Tulving 73] suggested his *encoding specificity theory* based on a distinction between episodic and semantic memories. The theory maintains that, "An encounter with a word results in the creation of a unique trace...." [Watkins 75], and implies that a subject does not learn, for example, words but experiences. This principle, 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. The emphasis is on two aspects: The existing retrieval cues and the characteristics of the past encoded episodes. What improves the retrieval is the degree of match between these two. The statement that,"Specific operations performed on what is perceived determine what is stored, and what is stored determines what retrieval cues are effective in providing access to what is stored" is accommodated also by experiments reported in [Barclay 74]. They showed how the word PIANO could be encoded differently when its different aspects are emphasized at the time of encoding. The subjects were given sentences such as "the man tuned the piano", or "the man lifted the piano". The cues "something melodious" and "something heavy" are used at retrieval time. The recall of piano was higher when the cue was "something melodious" and the learned sentence was "The man tuned the piano", and similarly when the cue was "something heavy" and the sentence was "The man lifted the piano". This implies that only specific aspects of piano was encoded each time, and the combination of the cue and the original encoding was necessary for the recall. Craik and Lockhart's *levels of processing hypothesis* [Craik 72] 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 to encode.

Realizing that Hewitt's context taxonomy fails to account for some of the empirical findings, and being inspired by the above-mentioned two principles, Baddeley suggested a distinction between *interactive* and *independent* context types. Baddeley's formulation is based on the role of context in learning, as well as in retrieval. He argued that interactive context "influences the memory trace directly by affecting the way in which the stimulus is encoded" [Baddeley 82], 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. "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 82].

Regarding the question of the relationship between recall and recognition, Baddeley favors a two-stage recall model where the candidates are accessed in the first stage. In the second stage, which he refers to as the recognition stage, the discrimination occurs [Baddeley 90]. Baddeley suggests that only interactive context elements affect the recognition stage while independent ones may also affect the candidate generation stage of recall. This may serve as a plausible explanation for why some environmental features affect verbal learning and face recognition, while others do not.

4 A model of context - a high-level ontology

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. We do not study problem solving separate from the problem solver, as has traditionally been done in AI. 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 it is occupied. Our interpretation of the findings from experiments described in section 3 highlights two important factors which help to recognize context typing. One is the *nature and demands of the process* that takes place during learning and remembering, the other is the *ground facts* that happen to exist in a situation.

We propose a distinction between *internal* and *external* context types in order to reflect these factors, where internal context relates to the former, and the external one reflects the latter factor. The key criteria for this distinction is the *deliberate* activity of subject's mind. The aspects regarding the agent is referred to as internal context. By this distinction, we emphasize the active role of the subject in fulfillment of various tasks. The type of cognitive process behind the reasoning is partly decided by the agent itself. The selection of the type of process can be imposed by giving instructions. This was the case in the experiments related to face recognition. The other alternative is that the subject himself decides how to process the material. This is particularly important in rather complex tasks such as problem solving, where the agent essentially replaces instruction with his own choice of cognitive behaviour. In our account, the agent's decisions regarding encoding details are shaped by his goals, hypotheses, and expectations, i.e. the internal context.

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. Our classification takes the *agent* (i.e. the reasoner) as the main criterion, while Hewitt’s distinction is based on the *target* as criterion. Hence, 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.

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.

5 Perspectives - the importance of internal-interactive contexts

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 91], [Mittal 93], [Pichert 77], [Lester 91]). For example, regarding text comprehension, Pichert and Anderson [Pichert 77] 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 ([Barclay 74], [McCoy 89], [Edmondson 93], [Cahour 93]).

including both non-contextual and contextual features and cues. A coherent set of aspects constitutes

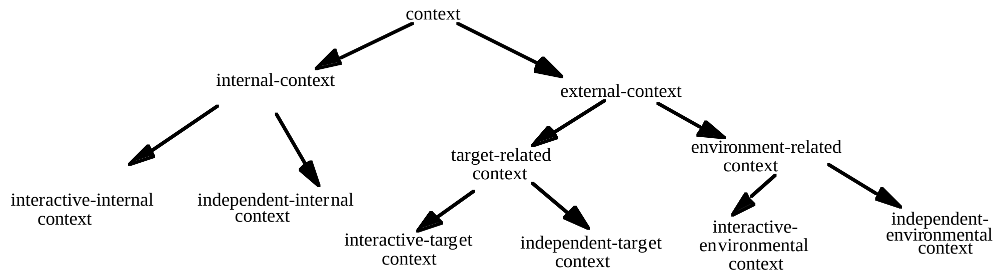

FIGURE 1. Context ontology emphasizing the active role of the reasoner

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.

Our account of the relation between the notions of context and perspective establishes a chain that starts at the goal. The goal determines the task necessary for its achievement. Once selected, the task then contributes to imposition of the perspective. The internal interactive context element, goal is thus a primary factor for identification of perspective. The goal of moving a piano, for example, makes the feature 'weight' relevant. It is commonly agreed that goals are of crucial importance for contextualized problem solving and learning (e.g. [Ram 91], [Leake 93], [Bogdan 94]). Similar context elements that contribute to an appropriate perspective are the predictions, expectations and background of the reasoner.

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 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.

6 Context in clinical problem solving

Two tasks in which a clinician extensively uses context elements are generating hypotheses about the cause of a failure, and planning actions. Hypothesis generation is often achieved by retrieving past episodes where similar features have been experienced [Tulving 73, Schank 82]. 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 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 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 hypothesis. For example, whether the patient has been operated for a fault in 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 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.

Figure 2 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 factors. For example, context cues such as age, sex, and 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 a domain model with an explicit representation of context. Domain knowledge that comprises the knowledge of facts regarding a particular domain is represented as three distinct, but yet closely connected components:'core-domain', 'external-context-domain´, and 'internal-context-domain'.

Figure 3 shows 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. 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.

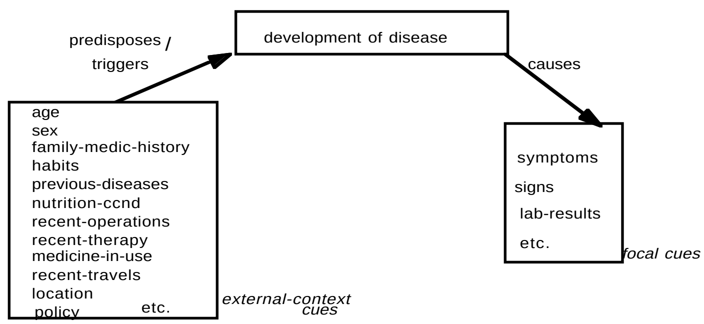

FIGURE 2: The Boundary between target-related context and non-context cues in medical domain

Context related domains also have a vocabulary and an ontology that facilitate an explicit representation of context-related concepts, and this, in turn, makes it possible to reason with and about context. In most cases, real world problems necessitate problem solving under highly uncertain conditions and with incomplete information.

Medical problem solving is an excellent example. An expert clinician always uses his past experience when encountering a new case. In this reasoning model, the use of 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. Throughout the whole process of problem solving, he gathers more non-context cues that enable him to solve the problem. In addition, he uses additional cues, namely contextual ones that are already available.

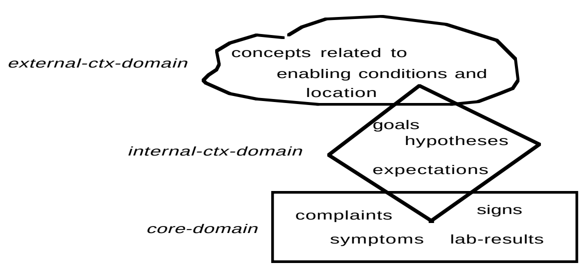

FIGURE 3. Domain knowledge augmented with context elements

As shown in figure 1, external context is grouped into target-related (i.e. patient context 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.

On the other hand, environmental information includes that related to equipment/tool availability, 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 expectations (consequences of hypotheses, expected results from actions, etc.)

7 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 (e.g. [Grimnes-96]). 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 Creek system [Aamodt 94a] 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 implemented in Lisp [Aamodt 94b].

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.

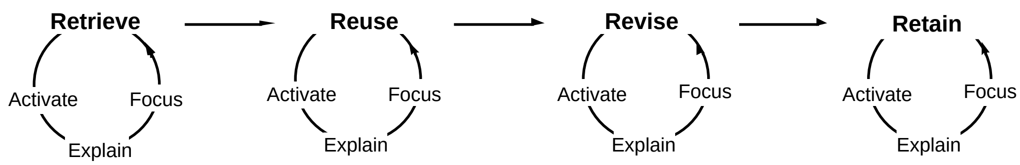

FIGURE 4. The CBR process and explanation engine

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: referred to as 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 diagnostic task structure is illustrated, related to an example run of the prototype implementation. The upper part of the figure shows the task hierarchy linked to some example cases used in the example run below. To achieve a diagnostic goal, two subtasks must be invoked: Generating a diagnostic hypothesis, and evaluating the hypothesis. As shown, 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 evaluation step.

Generation of hypotheses: The task of generating hypothesis involves retrieving a set of candidate cases from the case base, followed by selecting the best one. This case is then evaluated. The inputs to the task of generating candidate cases are the current goal, the new case, and the case base from which the candidate cases will be generated. The output is a set of candidate cases. The second subtask of hypothesis generation, selection of the best case, takes as input the set of candidates, the new-case, the general domain knowledge base, and knowledge related to the task. The output is the most similar case to the new case. In psychological terms, this type of retrieval is a “cued recall”. We adopt the two-stage recall model (see the last part of section 3) where the first stage generates candidates, and the second selects the best candidate. The second stage can be viewed as recognition.

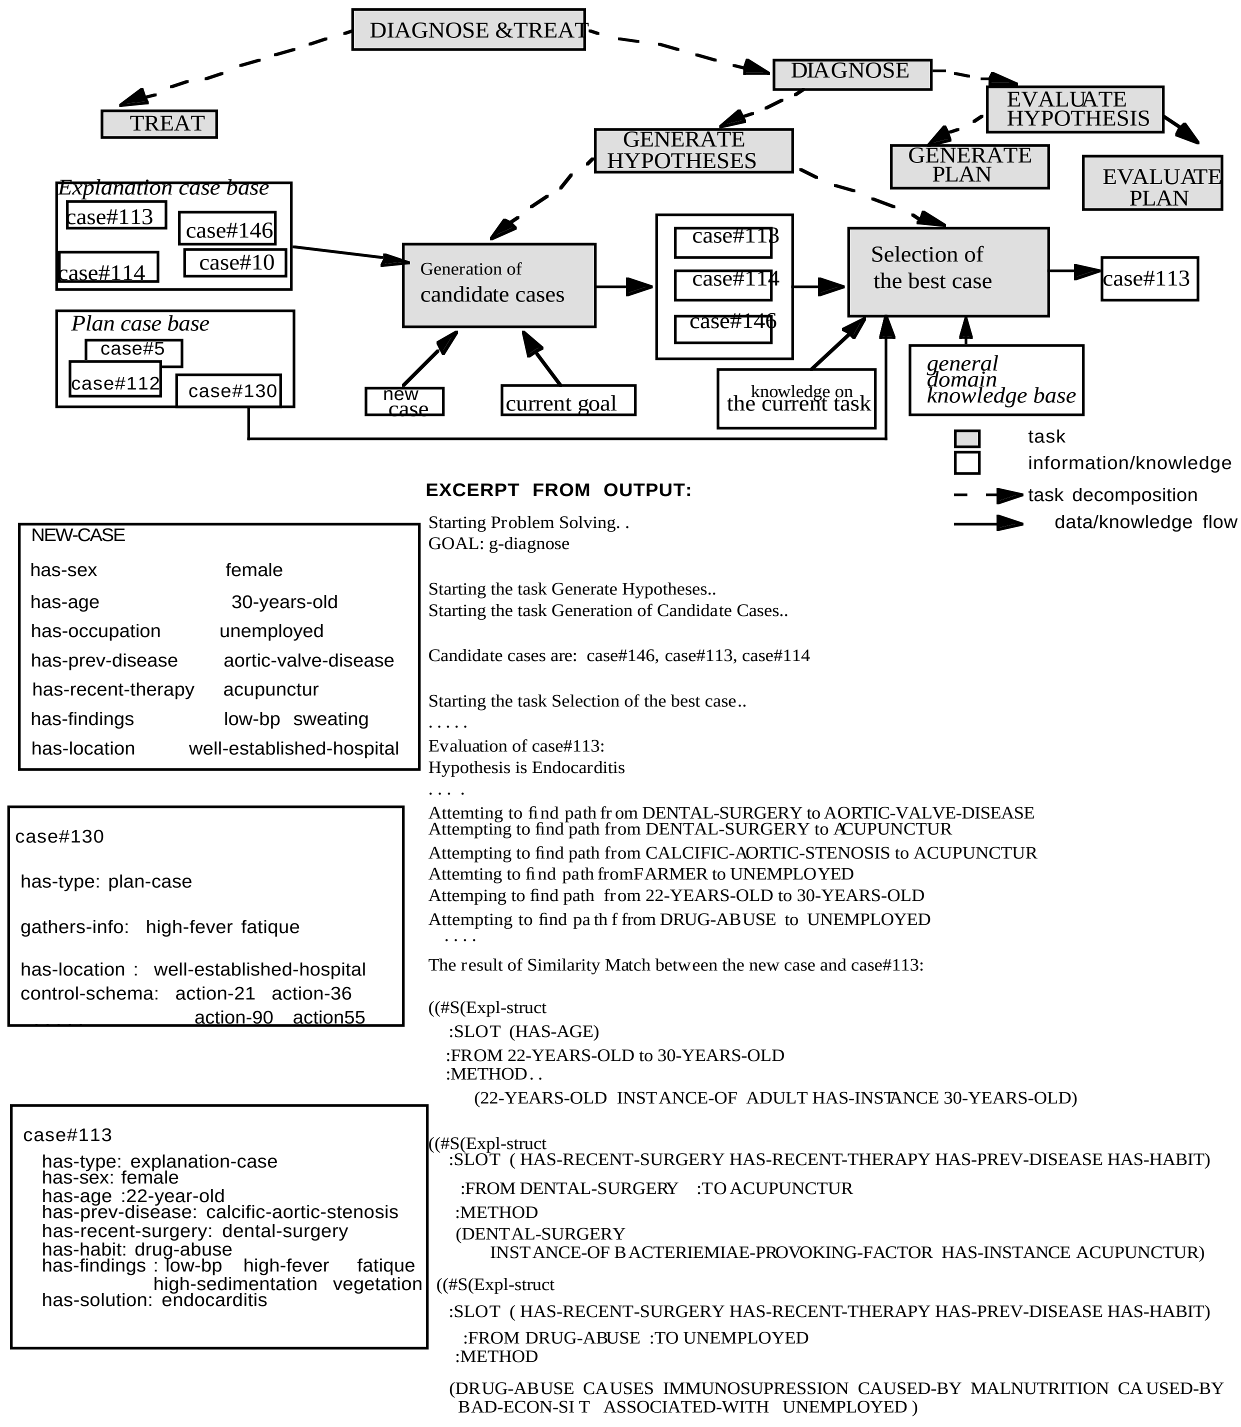

FIGURE 5: Diagnosis tasks, case structure, and excerpt from example run

in the new case are initially activated. The cases for which the total activation strength (combining all features) exceeds a threshold value are selected as candidates.The selection of the best case is based on the similarity judgment of candidate cases. Similarity 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. Indirect match happens when a path between slot values of the two cases is found in the general domain knowledge. For example, 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 contribute to the similarity of two 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 decision of how extensive this process should be influences efficiency.

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. 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. 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. 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.

Evaluation of hypotheses: 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 explanation case retrieval, we utilize interactive external context elements only, and not independent ones, despite that findings from empirical studies in cognitive science agreed 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. It is essential 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, for example, independent-target-context is ignored.

8 Summary and Conclusion

Most existing computational systems assume a single world that results in a decontextualized knowledge base. On the other hand, systems that attempt to take contextual effects into account often do so implicitly and arbitrarily. For example, in some rule based systems, context related knowledge is spread into rules, without any explicit way of reasoning with it. The existence of context in such systems (of which MYCIN was an early and simple example), is occasional rather than deliberate. So, in the reasoning process such context knowledge has no distinct role.

Because of the nature of the clinical diagnostic task, the retrieval cues do not sufficiently restrict the retrieval. This resembles the problem with ambiguous words. Selecting the relevant meaning of a word needs contextual assistance. It is the context that determines which interpretation is the relevant one. Similarly, given a limited set of features that may be observed in the existence of more than one disease, any context element that can help in narrowing the number of possible interpretations would be of crucial help.

Future knowledge-based systems need to become more flexible. The flexibility of a system can be ensured by explicitly representing various situations in which the system may apply, and allowing for partial matching of new problems to those situations. What is necessary for a flexible system? Two important factors are the capability to represent diverse circumstances in which a behaviour may emerge, and the ability to effectively utilize such knowledge. Context knowledge is utilized in choosing between alternative lines of reasoning. In this way, context serve a means for pruning the search in the problem space. In summary:

- It guides the search process in the solution space by imposing pragmatic biases, such as applicability. This is mainly achieved by use of external context. External context imposes constraints for choosing a method, a specific line of reasoning, which in turn leads to a more *relevant* and qualitatively better solution.
- It guides the process of finding a plausible solution. That is, it reduces the search for a plausible solution by applying a more dynamic and guided search. This is achieved mainly by the internal context. Internal context imposes perspective via the problem solving goal, which in turn leads to a more *focused* and efficient reasoning process.

The aim of the research reported here is to investigate how context can be taken into account in problem solving and learning, specifically diagnosis and treatment. In that sense, we share similar concerns with researchers who are interested in the ‘ practical’ aspects of context related to knowledge modeling and reasoning (e.g., [Abu-Hakima 95],and [Turner 94]). We feel, on the other hand, that studying the more theoretical and psychological aspects of the context is helpful in achieving this goal. The way we utilize contextual knowledge differs from [Abu-Hakima 95], as they are inspired from McCarthy’s axioms when defining context. Turner adopts a schema-based approach where contextual knowledge is represented in form of schemas. Our approach differs from Turner’s work primarily in that we use cases instead of schemas to capture episodic contextual knowledge. Maintaining unique episodes as cases, instead of generalizing them to schemas, does not abstract away details that may be useful for the problem solving process. Since it is left to the problem solving process to decide what to reuse of a past experience, and make the necessary generalizations on the basis of the actual problem situations instead of a predicted one, flexibility is gained.

The emphasis on problem solving goal as the major context determining factor is important, as is also described by other researchers. Ram and Hunter [Ram 91] emphasize the importance of “knowledge goals” that capture the reasoner’s specific desires to acquire or infer information. Leake and Ram [Leake 93] supports the view that goals have a focusing effect on learning, and influences what to learn from an experience. So far, we have concentrated on the problem solving aspects rather than learning aspects. Knowledge goals corresponds to what we refer as internal context. The influence of a goal in our model can be summarized as follows: The goal determines the task, which in turn specifies useful slot partitions. The slot partitions (subsets of relations), together with the current hypothesis, imposes a perspective on the domain knowledge.

Work currently in progress include context effects on case learning, specializing the modeling framework for the medical domain, and extending the experimental prototype.

9 References

[Aamodt 91] Aamodt, A., 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 94a] Aamodt, A., Explanation-driven case-based reasoning, In S. Wess, K. Althoff, M. Richter (eds.): Topics in Case-based reasoning. Springer Verlag, 1994. pp 274-288.

[Aamodt 94b]Aamodt, A., 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, 1994.pp 836-839.

[Abu-Hakima 95] Abu-Hakima, S. and Brezillon, P., Principles for the application of context in diagnostic problem solving, in *Working notes of the IJCAI-95 workshop on "Modelling context in knowledge representation and reasoning*", eds. Brezillon, P. and Abu-Hakima, S. , 1995

[Anderson 73] Anderson, J. R. and Bower, G. H. , *Human associative memory*, Winston, 1973

[Baddeley 82a] Baddeley, A., Depth of processing, context, and face recognition , in *Canadian Journal of psychology*, 1982, Vol. 36, No. 2, pp. 148-164.

[Baddeley 82b]Baddeley, A., Domains of recollection , in *Psychological Review*, 1982, Vol. 86, no.6, pp.708-729.

[Baddeley90] Baddeley, A. *Human memory,* Lawrence Erlbaum , 1990

[Barclay 74] Barclay, J.R., et.al., Comprehension and semantic flexibilty, in *Journal of Verbal learning and verbal behaviour*, vol 13, pp 471-481, 1974.

[Bogdan 94] Bogdan, R. J., *Grounds for Cognition: How goal-guided behaviour shapes the mind*, Lawrence Erlbaum associates, 1994.

[Brezillon 93] Brezillon, P. , *Proceedings of the IJCAI-93 Workshop on "Using knowledge in its Context"*, Laforia

[Cahour 93] Cahour, B. and Karsenty, L., *Context of dialogue: A cognitive point of view,* in [Brezillon 93] , pp 20-29

[Craik72] Craik, F. I. M. and Lockhart, R. S. , Levels of Processing: A framework for memory research ,in *Journal of verbal learning and verbal behaviour*,1972, Vol.11, pp.671-684.

[Dalton93] Dalton, P., The role of stimulus familiarity in context-dependent recognition, in *Memory and Cognition*, 1993, Vol. 21 (2), pp. 223-234.

[Edmondson 93] Edmonson, W. H., and Meech, J. F., *A model for Human-Computer Interaction*, in [Brezillon 93], pp. 31-38.

[Eich 80] Eich, J. E., The cue-dependent nature of state-dependent retrieval, in *Memory & Cognition*, 1980, Vol.8, pp. 157-173.

[Godden 75] Godden, D.R. and Baddeley, A. D., Context dependent memory in two natural environments: On land and underwater. , in *British Journal of psychology*, 1975, Vol. 66, pp. 325-331.

[Grimnes 96] Grimnes, M., Aamodt, A., A two layer case-based reasoning architecture for medical image understanding, in *Proccedings of EWCBR*, 1996

[Kintsch 74] Kintsch, W. The represenation of meaning in memory, Lawrence and Erlbaum, 1974

[Leake 93] Leake, D. and Ashwin R., Goal-driven learning: Fundamental Issues and Symposium report, in *AI Magazine*, winter, 1993.

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

[McCoy 89] McCoy, K. F., Generating context-sensitive responses to object-related misconceptions, in *AI Journal* , vol 41, pp 157-195, 1989.

[Mittal 93] Mittal, V. O., and Paris, C. L., Context: Identfying its elements from the communication point of view, in [Brezillon 93], p 87-97.

[Pichert 77] Pichert, J. W., & Anderson, R. C., Taking different Perspectives on a story, in *Journal of Educational Psychology,* 1977, Vol.69, No. 4, 309-315.

[Ram 91] Ram, A. and Hunter, L., A goal based approach to intelligent Information retrieval, in Machine Learning: Proceedings of the 8. International Conference, 1991, pp 265-269

[Schank 82] Schank, Dynamic memory: A theory of learning in computers and people. Cambridge Univ. Press..

[Suthers 91] Suthers, D. D., A task-apropriate hybrid architecture for explanation, in *Computational Intelligence*, vol.7,1991.

[Tiberghian 86] Tiberghian, G. Context and cognition: Introduction, in *Cahiers de psychologie cognitive. European bulletin of cognitive psychology*, 1986, Vol.6, No. 2, pp. 105-119.

[Thomson 88] Thomson, D. M. and Davies, G. M. *Memory in context: Context in memory*, Wiley, 1988.

[Tulving 73] Tulving, E. and Thompson, D. M., Encoding specificity and retrieval processes in episodic memory., *in Psychological Review*, 1973 , Vol. 80, pp. 353-370.

[Turner 94] Turner, R., Adaptive reasoning for real-world problems, Lawrence Erlbau associates, 1994

[Watkins 75] Watkins, M. J., and Tulving, E., Episodic memory: When recognition fails, in *Journal of experimental psychology:General*, 1975, vol. 104., No.1, 5-29.