Knowledge-Level Modularization of a Complex Knowledge Base

Josef Meyer-Fujara, Barbara Heller, Sonja Schlegelmilch, Ipke Wachsmuth

Faculty of Technology, University of Bielefeld, D-33501 Bielefeld

Abstract. An increasingly relevant theme for knowledge-based systems (KBS) is how to model and use a large and complex knowledge domain. On the one hand, this involves developing new ideas on a modular construction of a knowledge base and, on the other hand, an appropriate architecture of a KBS that takes account of such a modular knowledge base. In the HYPERCON project[1] we conceptualize a complex medical domain (hypertension) according to ideas of a knowledge-level modularization. In this paper, we give a brief account of our approach and describe how specific knowledge is structured and accomodated in a modular knowledge base. Using a specific case of a patient, we exemplify focussing procedures based on plan-and-tactics and on changing between different granularity levels in the course of a consultation.

1 Introduction

One of the severe limits for the application of knowledge-based systems still lies in the restricted size of manageable knowledge bases. Many fields of application, however, require the handling of complex – i.e., large and diverse – stocks of knowledge. This issue has been noticed for some time, e.g., [18, 10, 2, 21, 16, 13], and its actuality is witnessed by a growing number of activities at major AI conferences.

Complex domain knowledge can neither be acquired nor implemented and maintained by a single person. But division of labor necessitates that team members can restrict their attention to limited parts of the overall knowledge. Efficiency requires that the problem solver strictly focus on parts of knowledge relevant for the actual issue if the system is not to drown in memory search. Hence, a modularization of KBS seems necessary. This task cannot be addressed solely in a way that is convenient for software development. In our view, modularization must follow semantic borders of relevance and must be addressed at the *knowledge level* in the sense of Newell [14]. Such a modularization should entail three questions:

1. Which criteria guide modularization?
2. How is relevant knowledge actually retrieved?
3. How is access to the relevant knowledge organized?

---

1 Our work is supported in part by the North-Rhine-Westphalian Ministry of Research and Technology (MWF grant IVA6-400 015 92). We thank the medical experts Prof. K. Kauffmann, Prof. H.-D. Faulhaber and Dr. U. Müller-Kolck for their patience in explaining medical facts as well as our students Ch. Dücker, J. Hamann, A. Möller, Ch. Scheering and J. Stoye for their assistance.

We deal with these questions in the context of the HYPERCON project which aims at supporting diagnosis and therapy of hypertension patients [20]. The interest of the system lies in helping the user to compare possible diagnostic procedures, diagnoses and therapies in the light of heuristic and deep knowledge automatically related and applied to his actual case.

A very large body of knowledge has to be represented pertaining to diverse fields of medicine such as laboratory tests, image generating procedures, epidemiology, and pathophysiological models. In addition, there is knowledge about reasonable courses of actions for finding a diagnosis that reduce strain and cost to the patient.

The knowledge is of diverse origin (from textbooks, expert interviews, case studies) and it is needed at different levels of detail: For dealing with standard cases, compiled instructions for action are sufficient. For taking care of, e.g., multimorbidity, however, and also for explanation, detailed knowledge is needed. Morphological knowledge is needed on size levels reaching from organs down to electron-microscopical structures.

Finally, the particular parts of knowledge may strongly interact: a heart murmur may be caused by high blood pressure originating from a kidney damage caused by diabetes, but may as well be a symptom of a primary heart disease. In total, our domain clearly illustrates the need for modularization.

In our current work, we concentrate on the highly relevant subfield of hypertension due to kidney diseases or hormonal dysfunctions, possibly complicated by diabetes. Our system is realized on top of the hybrid expert system shell Knowledge Craft.

After discussing related work, we will explain our approach to modularization with respect to knowledge partitioning, knowledge focussing, aspects of granularity, and system architecture. Then we will illustrate the concepts introduced by the example of an authentic medical case.

2 Previous Approaches

Known approaches to modularization of knowledge representations (as opposed to modularization of knowledge-based *systems*) have concentrated on several aspects. Since long, possibilities of syntactical rule grouping have been available, e.g., in variants of OPS such as RIME [18], and in major expert system tools. For frames, such tools often offer so-called contexts or worlds to restrict visibility of slots and values. On a more knowledge-related level, Davis [6] already suggested content-directed invocation of rules by means of meta-rules. Clancey [4] has proposed to differentiate knowledge according to its use (e.g., for problem solving or explanation), and to separate control knowledge from domain knowledge.

This latter separation is favored and refined in KADS [22] where several layers of control knowledge are distinguished in the acquisition-oriented *conceptual* model. Soloway et al. [18] suggest domain specific "buckets" (hierarchically organized *problem spaces*) that correspond to common purposes such as a variety of configuration functions. Purpose and content ("topic") are emphasized also by Clare [5]. In the "Knowledge Sharing Effort" – cf. [13, 15, 7] – a possibility to tackle the problem of realizing large knowledge bases is seen in the reuse of already existing knowledge bases; accordingly, efforts for standardization and for the establishment of libraries are made.

Finally, starting from the observation that any particular axiom (e.g., rule) will be irrelevant in many contexts, the incremental acquisition and representation of self-contained clusters of domain knowledge in a *partitioned knowledge base with dynamic access conditions* have been suggested. Their basic principles consider content and specificity of knowledge as structuring aspects and suggest to organize knowledge in layered, possibly overlapping *knowledge packets* [19].

3 Modules, Knowledge Packets and their Focussing

Our approach is led by the following motivations: the knowledge base should be partitioned into parts

- that are competent for definite task domains
- that may hold specific knowledge representations and control strategies
- that may hide their contents from other parts
- that, to a large extent, may be developed independently
- to which changes, debugging and consistency checks may be restricted
- that are manageable in themselves and allow the aggregation of more complex knowledge parts.

Our examination of the domain suggests to differentiate knowledge to be modularized by the following criteria:

- simultaneous use in the problem solving process
- cohesion of content
- similar granularity level (size, abstraction)
- specificity (applicability in particular circumstances, usability for particular goals).

knowledge about	knowledge about	criteria
resting electrocardiogram	exercise electrocardiogram	cohesion of content
findings that give hints at a disease	findings that confirm a disease	simultaneous use
macroscopic parts (e.g. renal calices)	microscopic parts (e.g. glomeruli)	granularity (size)
comparison of resting/exercise/ long-term electrocardiogram	interpretation of resting/exercise/ long-term electrocardiogram	specificity (usability)
physiological parameters during pregnancy	normal physiological parameters	specificity (applicability)

Fig. 1. Separable knowledge parts (first and second column) according to differentiating criterion (third column).

For illustration, fig. 1 shows a table of knowledge parts that will be kept separate from each other, and indicates the relevant criterion. In the rest of the paper, further examples will be found.

3.1 Knowledge Modules and Knowledge Packets

By intensive discussion of diverse authentic cases with medical experts of different background (theorists, clinicists, and a practitioner), a general model of the diagnostic reasoning process was established (cf. fig. 2). The overall diagnostic plan suggested by this model includes different stages that ideally follow one another, although in reality often are subject to iteration. These processing steps were found to concern patient history, physical examination, laboratory tests, diagnostic procedures, hypothesis generation, hypothesis verification, and therapy. Correspondingly, *knowledge modules* were defined that are conceived to be highly independent and active one at a time. In a less chronologically separable way, knowledge about nosology and about physiological and anatomical models was found to be used.

All knowledge modules are characterized by self-contained specific knowledge and well-defined interfaces towards a coordination component. The interfaces hide the interior representation choices and the specific knowledge.

Altogether, these are the following modules:

acquisition modules:	patient history, physical examination, laboratory tests, diagnostic procedures
diagnostic modules:	hypothesis generation, hypothesis verification
therapy module:	drug treatment, non-drug treatment, invasive treatment, operative treatment
library modules:	nosology, anatomical and pathophysiological models

The knowledge modules are subdivided into *knowledge packets*, assembling collections of knowledge elements to be focussed simultaneously.
The packets may properly contain further packets as particular bodies of knowledge may branch to extend in competitive subbodies. Knowledge elements not in focus are invisible to the inference engine in order to exclude irrelevant knowledge from searching and matching. The set of packets is structured hierarchically according to their degree of specificity. For illustration, an extract of the packet structure "image generating procedures" in the module "diagnostic procedures" is shown in fig. 3. The overall knowledge base has approximately 250 knowledge packets.

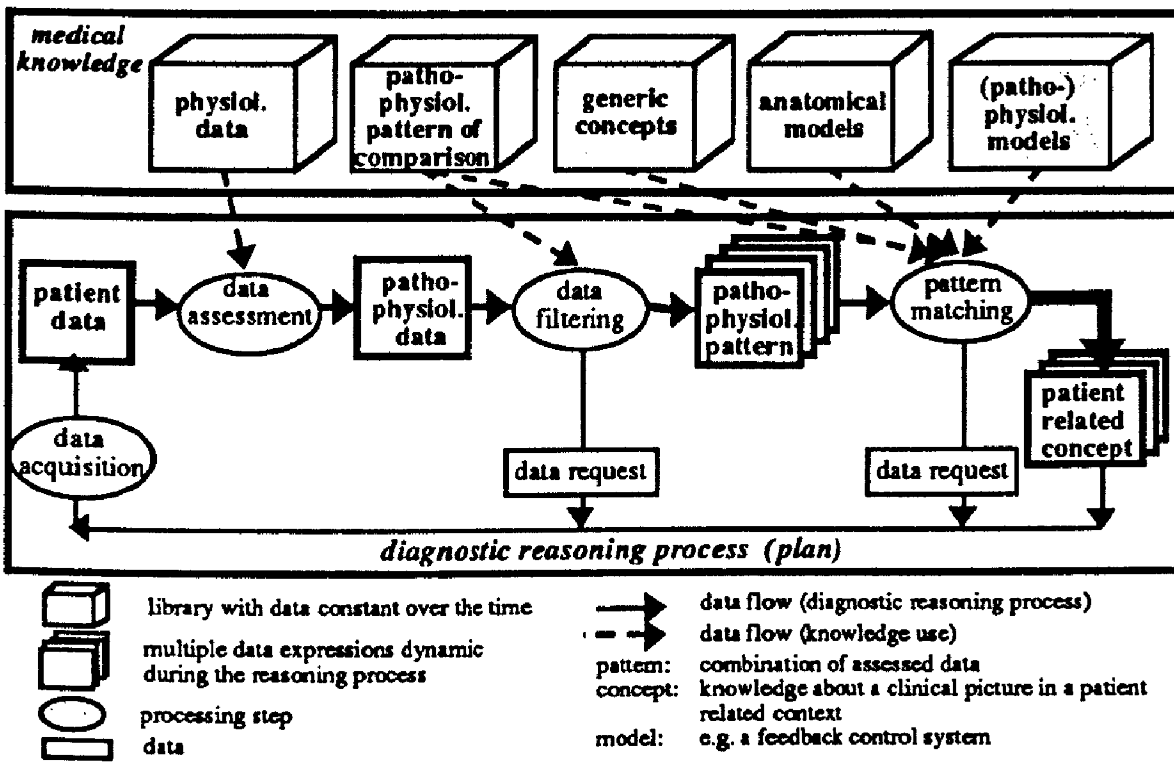

Fig. 2. Model of the diagnostic reasoning process (further explanation in text).

Packets shown side by side are competitive with each other. Competitive packets may represent alternative views or conflicting knowledge. Besides domain knowledge, modules (always) and packets (often) must contain *control knowledge*, describing, e.g., in which circumstances subpackets are focussed.

Similar to KADS the entire knowledge is organized in five levels, namely, strategic, task, tactics, inference, and domain level. This allows for smooth extension of packets in case the KBS must be enlarged for handling further tasks. Control knowledge, in our opinion, must be split along with domain knowledge. E.g., unless image generation procedures are considered at all, the choice among these procedures is irrelevant. Also, whether to proceed by establish-refine or generate-and-test depends on the local disease heterarchy. So packets have their own five-level control knowledge. In smaller packets, the three upper levels may be unimportant. We will explain the levels in section 3.3.

3.2 Partitioning Knowledge According to Granularity

In our opinion, *granularity* [9,3] must be taken into account as a special aspect of cohesion [23]. Reasoning most often occurs at a definite granularity level that is switched if necessary, and primarily to an adjacent level. Thus knowledge packets should contain knowledge at a comparable grain size. Two aspects of granularity seem particularly important for knowledge modularization.
Foundation concerns the bottom level of hierarchies used, that pertains to the objects taken for atomic, i.e., needing no further differentiation. For instance, the foundation of anatomic knowledge needed for the interpretation of sonographic findings about the kidney lies at macroscopic structures such as renal pelvis and artery; due to the restricted resolution of sonography, smaller ones (e.g., glomeruli) need not be considered. Only when this restriction turns out to be too strict, other levels are called in. Analogously, a finding of proteinuria gives a hint at a kidney disease as opposed to a heart disease, but with this finding alone, knowledge about differentiating kinds of specific kidney diseases are not relevant.

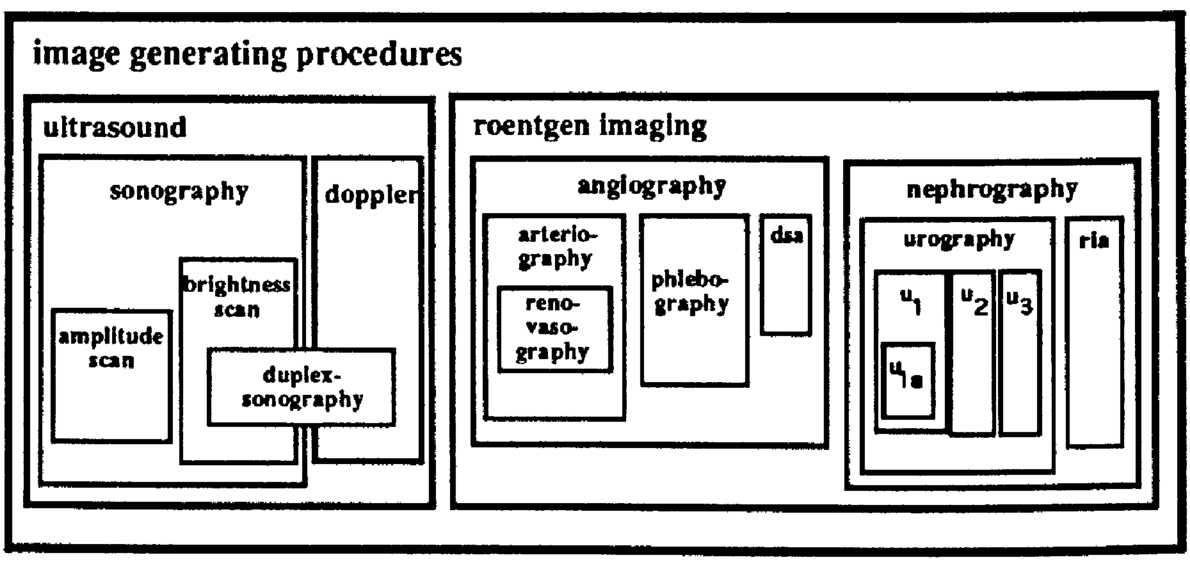

dsa digital-substraction-angiography; ria radio-isotope-angiography; u_1 extraction urography; u_{1a} micturating urography; u_2 infusion urography; u_3 retrograde urography

Fig. 3. Extract of the packet structure inside the module "Diagnostic Procedures".

The second aspect of granularity is the *field of view*, i.e., the region of physical or abstract entities that are focussed at the same time. It is a special kind of context as formalized, e.g., by McCarthy [11] and determines the interpretation of names of entities (e.g., leucocyte-count) and inferences associated with them (e.g., assessment of normality). As far as possible, our knowledge packets admit a well defined field of view to be automatically used when the packet is focussed. Technically, packets can be associated with Knowledge Craft `contexts` and contain context-dependent schemata and rules.

Defining packets according to granularity entails the need to link specific terms to more general ones at changes of packet focus. For example, the packet P "therapy of chronic kidney diseases" suggests protein-reduced dieting. A diagnosis of chronic-pyelonephritis cannot lead to this suggestion unless it is classified as chronic-kidney-disease. But this classification knowledge is contained neither in packet P nor in more general ones: P has incomplete knowledge about the hierarchy its entities belong to.

This dilemma is tackled by packet interfaces: By them, the diagnosis chronic-pyelo-nephritis suggested is linked there to the taxonomic hierarchy. The interfaces also deal with partonomic foundation and with view-dependent terms.

3.3 Global and Local Focussing of Knowledge

Focussing of knowledge is done at two abstraction levels: choosing among modules (*global focussing*) and choosing among packets inside a module (*local focussing*).

Global focussing (GF). During problem solving, the module to be activated next is determined by the global focussing component. It is structured in a way similar to KADS [22] (cf. fig. 4).
On the *strategic level*, in our case, the problem class is preset, succcessively to "diagnosis" and to "therapy", and the problem-solving method is preset, in our subsequent example, to "heuristic classification". Efficient focussing of knowledge is based on global strategies, which are described by global plans and different alternatives within the plans. They describe a global procedure of inferencing which corresponds to abstractions of our model of the diagnostic reasoning process shown in fig. 2.

Strategic Level	problem classes problem-solving methods global strategies global plans
Task Level	classes of tasks, partial tasks, subtasks
Tactics Level	default plans alternative plans competitive plans
Inference Level	rule-classes for description of inference steps during problem-solving process, rule-classes for plans
Domain Level	modules

Fig. 4. The levels of global focussing.

The *task level* describes the particular tasks for diagnosis (e.g., establish hypothesis) which are structured in different subtasks. These can be instantiated one at a time.

To obtain a more specific focussing within the modular knowledge base, we introduced a further level – *the tactics level* – as an extension of the KADS-model. Tactics describe methods for situation-dependent application of inference steps according to the goal of the actual subtask. Default tactics may be abandoned for alternative tactics. E.g. the default tactics for generating a hypothesis is "patient-related". If the system cannot generate a hypothesis from the existing patient data, the alternative tactics of epidemiology-based hypothesis generation is used.

In the REFLECT-project [17], a framework with a reactive planning task that makes strategic decisions about optimal sequences of problem solving actions is given. In contrast, in our approach, the tactics level serves as an integrated link between task and inference level and allows a situation-based preparation and use of problem solving actions (e.g. by loading and activating rule-classes).

At the *inference level*, rule-classes implementing inference steps and tactics are to be found.

At the bottom level (*domain level*) of the GF, the modules and their competence are described.

In order to further focus domain-specific knowledge inside a module being focussed on the global level, a suitable interface must exist between the GF and the local focussing component (LF) of a packet. The interface is described by the following contents:

- global situation (consisting of the task, subtask and the suggested tactics)
- patient related concept (consisting of processed patient data and already obtained results, e.g., patterns, hypotheses, ...)
- chosen module.

Local focussing (LF). The local focussing component focusses part of the domain knowledge inside a module or a knowledge packet. In the following, for simplicity we will assume to be dealing with a module. There is a specific LF for each module. It uses the module's control knowledge and is subdivided into five levels similar to those of the GF. The local strategy and the task level of the LF correspond to a refinement of levels of the global focussing component depending on the particular content competence of a module. This means: If the GF focusses, e.g., the module "Diagnostic Procedures", the global subtask "Generate Hypothesis" becomes a local strategy in the LF of the module. Depending on this strategy and on the local plan, now the tasks to be executed by this module are determined. According to the associated task-specific tactics, the corresponding knowledge packets in this module are focussed and activated.

The contents of the different levels in the LF of the module "Diagnostic Procedures" look as in fig. 5 for the medical case described in section 4. There the reader will also find a more precise trace of focussing.

The domain level consists of descriptions of packets with respect to their specificity, degree of abstraction, etc. Normally, the LF chooses a packet just one step down the packet hierarchy.

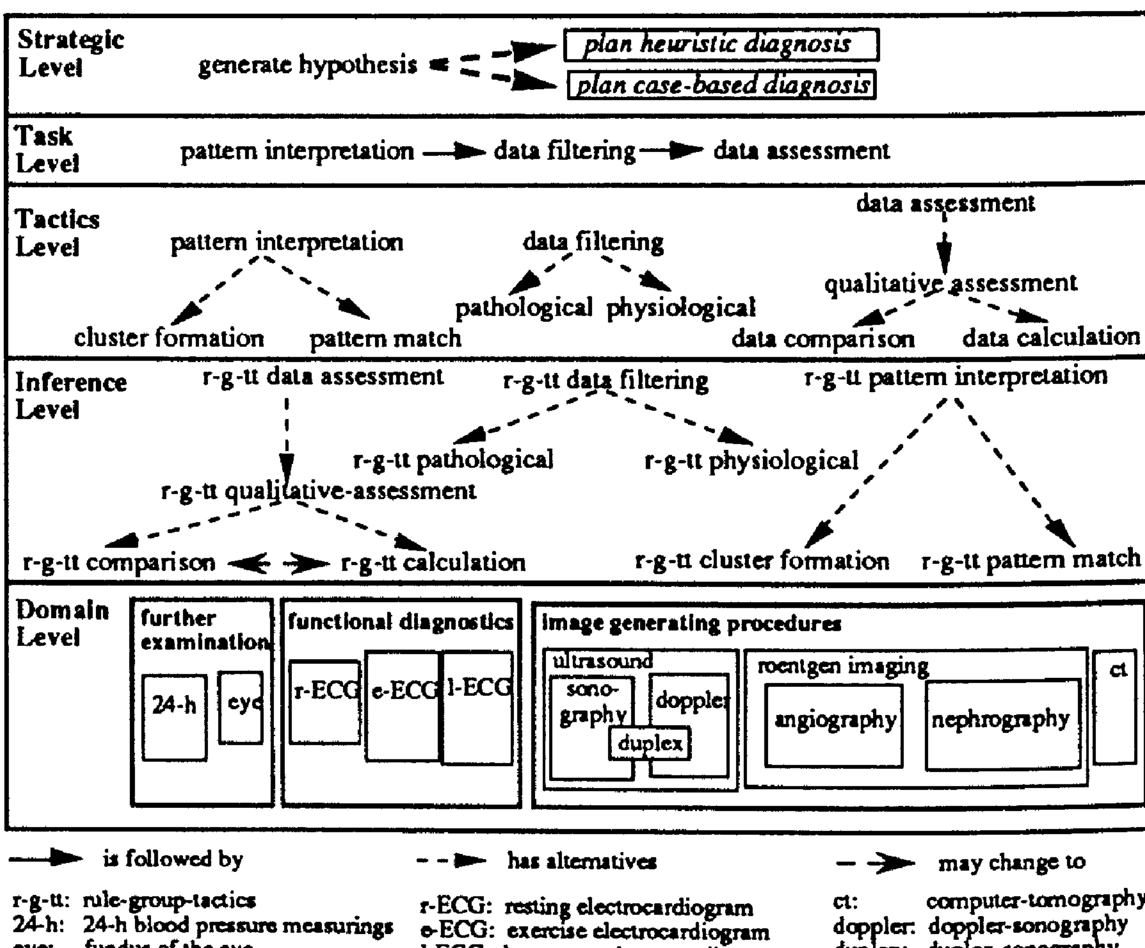

Fig. 5. LF of the module "Diagnostic procedures". On the domain level, for lack of space, only the packets needed to diagnose kidney artery stenosis are shown.

Rules may be focussed according to packets they belong to (static structure), but also according to object classes and attribute names occuring in the rules' premises or conclusions (dynamic structure). This possibility was created by associating a description schema to each rule.

3.4 System Architecture

The most important components of the system architecture are the global focussing component (GF), the knowledge base component (KB) and the coordination component (CO) (cf. fig. 6).

The global focussing component discussed above manages the knowledge-related interaction of modules. It carries out the overall plan by successively activating modules, examining their results and continuing or modifying the plan accordingly. The knowledge base component consists of the knowledge modules as described in section 3.1, where each module possesses its own local focussing component. Some of the modules are shown in fig. 6.

The coordination component organizes the data transfer between the user and the system. It contains descriptions of the actual situation and data and is responsible for furnishing them to and obtaining them from the GF and thus to and from the modules. Each transfer process is managed by a standardized communication structure [8, 12]. Thereby the independence of the knowledge-related components GF and KB from technical specificities is supported.

In the following section the diagnostic reasoning process will be explained on the basis of a patient case.

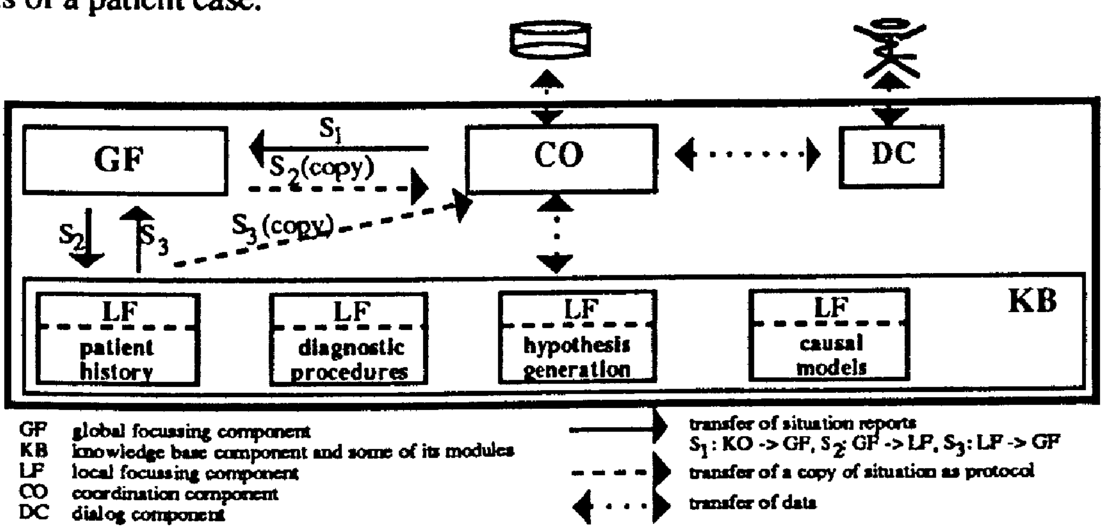

Fig. 6. Architecture for a knowledge-based system with modular knowledge bases.

4 Example

Now we present a trace of the focussing procedure during processing an authentic case (obtained from one of our medical experts). The following background is given: A middle-aged woman consults a physician (a general practitioner) who has not met her before. From the physician's input concerning the patient data an adequate pathological patient concept is constructed. (Geneva font indicates system output to the user).

Pathological patient concept:

History: female, 52 years old, insulin-dependent diabetes mellitus (IDDM) and frequent pyelitis in the past.

Physical exam.: 160 cm, 85 kg, blood pressure 160/90 mmHg.

Laboratory test: blood leucocyte count 13000/ml, haemoglobin 11 g/dl, glucose 200 mg/dl, blood-urea concentration 60 mg/dl, serum uric acid 8.2 mg/dl, creatinemia 3.6 mg/dl, creatinine clearance 25 ml/min, urine leucocyte count 75.

This concept is transferred to the GF.

(1) GF The GF must decide which module is suited best for the given status of the patient. The levels of the GF will be described by the following schema given that there is no hypothesis available yet:

global strategy: *<heuristic classification>*
global plan: *<default plan>*
main task: *<specify diagnosis>*
subtask: *<generate hypothesis>*

According to the ordinary plan and the subtask, the GF focusses the first module ("Patient History") and sends the patient history data to this module as part of the patient concept.

(2) LF The LF interprets the data obtained by the GF and this results in the LF having (cf. fig. 5):

local strategy: *<generate hypothesis>*
local plan: *<data assessment, data filtering, pattern interpretation>*

Considering the present task, the corresponding tactics and the related rule classes will be instantiated and activated. This leads to a condensation of information resulting in a more detailed focussing on specific packets inside the module. At first the patient history will be assessed. Depending on the underlying pathological and physiological knowledge about the patient, the obtained data will be filtered. Regarding the filtered data, a pathological pattern will be interpreted. The particular tasks will be processed step by step and the preliminary results will be transferred to the GF.

Output of the results of the module "Patient History":
> female, 52 years, insulin-dependent diabetes mellitus and frequent pyelitis
> explanation: possibly there is a dependency between diabetes and frequent
> pyelitis. Please check the status of the kidney.

The preliminary results of the module will be worked up and transferred to the global focussing component. Step (1) and (2) will be repeated for the modules "Physical Examination" and "Laboratory Test". For conciseness, for them we only describe focussing and results.

(3) LF The LF focusses the knowledge packets "Blood Pressure Measurements", "Basic Parameters", and "Vegetative Symptoms" inside the module "Physical Examination".

Output of the results of the module "Physical Examination":
> overweight, hypertension

(4) LF The LF focusses the knowledge packets "Serum Laboratory Tests", "Blood Count", "Urine Laboratory Test", and "Urine Sediment" inside the module "Laboratory Tests".

Output of the results of the module "Laboratory Test":
> creatinemia greatly increased, creatinine clearance decreased,
> blood-urea increased, serum uric acid increased,
> blood leucocyte count greatly increased, haemoglobin decreased,
> glucose greatly increased, urine leucocyte count increased.
> hypothesis: suspicion of (non acute) kidney disease.
> explanation: suspicion of kidney disease because of greatly increased
> creatinemia,
> decreased creatinine clearance and a slight proteinaemia; non acute
> because there are no indications of acute symptoms.

What to do now? This hypothesis is to vague to serve as a diagnosis and new patient data are not available. Therefore an alternative way should be looked for in the strategy plan of the GF. The plan provides consulting the module "Diagnostic Procedures" first.

(5) GF The GF focusses the module "Diagnostic Procedures" and sends it the relevant patient data and the inferred hypothesis.

(6) LF For reasons given below, the LF focusses the relevant part of knowledge inside the module "Diagnostic Procedures" and inside the "Anatomical Models" module.

By means of the present vague hypothesis of a kidney disease, the module suggests a kidney sonography as a first search test. To get better results in case of a possible kidney stenosis, the more specific packet "Brightness Scan" is chosen. For the interpretation of the findings, morphological knowledge is necessary. The relevant part of the library module "Anatomical Models" is determined by granularity considerations: The field of view is restricted to the kidney and its neighbors.
Except for the kidney, the foundation level lies at organs as such. For the kidney itself, it is calculated from spatial and tissue resolution of sonography. Thus knowledge about, e.g., normal size and thickness of kidney, pelvis, cortex, and marrow is focussed, but none about microscopic structures.

User-input The sonographical findings show a cirrhosis of the right kidney but no engorged kidney.

Using the morphological knowledge, the degree of kidney damage is deduced from observed kidney sizes.

(7) LF The results of the sonographical findings are transferred to the querying module. The module focusses first the packet "Pattern Interpretation" according to the global task "Hypothesis Generation". At last a pattern regarding the new patient data is built.

Output of the results of the module "Diagnostic Procedure":
> manifest kidney damage in form of a cirrhosis of the right kidney
> but no engorged kidney.
> pattern: cirrhosis of the right kidney, creatinemia greatly increased,
> frequent pyelitis, urine leucocyte count increased.

(8) GF Together with the relevant patient data the interpreted pattern will be transferred to the focussed module "Diagnosis Generation".

(9) LF The focussed packet which considers data from "Patient History", "Laboratory Test" and "Diagnostic Procedure" together attempts to generate a hypothesis considering the built pattern.

Output of the results of the module "Hyothesis Generation":
> diagnosis: chronical pyelonephritis.
> explanation: pyelonephritis: because of frequent pyelitis in the past;
> chronical because of cirrhosis of the right kidney ......

Before the inferred hypothesis will be displayed to the user, it is verified by the module "Hypothesis Verification". This module contains knowledge about competitive nosographies. After verification of the inferred hypothesis, the global plan for generating a hypothesis is exhausted and a new global plan for generating a therapy will be instantiated.

5 Summary and Outlook

In this paper, we explained our principles for a knowledge-level modularization of a complex (medical) knowledge base and, in particular, showed how their application permits to focus on restricted parts of knowledge.

Based on a model of the overall reasoning process, the knowledge base is decomposed into knowledge modules highly independent from each other, active one at a time. These modules are subdivided into knowledge packets according to specificity and granularity. Focussing such a packet makes knowledge elements outside of it invisible to the inference engine. Whereas *global* focussing activates a module, *local* focussing narrows focus further down to packets or subpackets. Both modules and packets contain associated control knowledge subdivided in KADS-like layers and linked by a layer shift during local focussing. Independent representation and granularity choices in different parts of the knowledge base are made possible by packet interfaces, a common interface handler and a central communication control component. The interplay of knowledge partition, dynamic focussing and the architecture components was illustrated by an authentic medical case.

The system is not in actual use yet; testing covers the entire span of renal hypertension. It relies on the users' competence for weighing evidence; we believe this issue to be rather independent of knowledge structuring.

A problem still consists in determining the useful size of packets. The main criterion is their cohesion: further subdivision should either seem unnatural or lead to crumbs not worth focussing. Sometimes, concurrent organization principles are conceivable, e.g., knowledge about blood and urine parameters may be grouped according to medium or else according to disturbances they give hints at (leading to groups such as kidney profile that contain parameters of blood as well as of urine). In this case, our approach has been to use the "simple" organization according to medium for the basic knowledge (usual physiological range of the parameters) and to keep knowledge about the kidney profile synopsis in a further packet. For optimal packet organization, careful analysis of the domain and of possible interactions is necessary. Sometimes more specific knowledge should be taken into account before more general knowledge has been tried unsuccessfully, especially when special cases admit particularly simple solutions. Criteria for removal seem of minor importance for the moment as the main benefit lies in activating relevant knowledge only.

To conclude, we have found the described approach of knowledge-level modularization greatly useful, especially with respect to an incremental knowledge base development. Though we have worked out modularization principles in the context of a medical domain, we expect their application to yield equally good or better results in other domains, especially when a domain is less strongly interconnected than our domain of hypertension.

References

1. Antoniou, G. & Wachsmuth, I.: Structuring and modules for knowledge bases: motivation for a new model. Knowledge-Based Systems 7 (1), 49-51 (1994)
2. Bocionek, S.: Modulare Regelprogrammierung. Braunschweig: Vieweg (1990)
3. Bylander,T. & Chandrasekaran, B.: Generic Tasks for Knowledge-Based Reasoning: the "Right" Level of Abstraction for Knowledge Acquisition. International Journal of Man-Machine Studies 26, 231-243 (1987)
4. Clancey, W.J.: The Epistemology of a Rule-Based Expert System – a Framework for Explanation. Artificial Intelligence 20 , 215-251 (1983)
5. Clare, M.: When Information Engineering Meets Knowledge Management. Proceedings AAAI Workshop Modelling in the Large, Washington (1993)

6. Davis, R.: Generalized Procedure Calling and Content-Directed Invocation. SIGPLAN Notices 12, 8, 45-54 (1977)
7. Ginsberg, M.L.: Knowledge Interchange Format: The KIF of Death. AI Magazine 12 (3), 57-63 (1991)
8. Heller, B. & Schlegelmilch, S.: Modularization of Knowledge: A Competence-Oriented Approach. Proc. AAAI Workshop Modelling in the Large, Washington (1993)
9. Hobbs, J.R.: "Granularity," IJCAI-85, pp. 432-435 (1985)
10. Lenat, D.B. & Guha, R.V.:. Cyc: a midterm report. AI magazine 11 (3), 32-59 (1990)
11. McCarthy, J.: Notes on Formalizing Context. IJCAI-93, pp. 555-560 (1993)
12. Heller, B., Meyer-Fujara, J. & Schlegelmilch, S.: HYPERCON: Ein Konsultationssystem zur Hypertonie auf Basis modular organisierter Wissensbestände. KI-94 Anwenderkongreß (1994)
13. Neches, R. et al.: Enabling technology for knowledge sharing. AI Magazine 12 (3), 37-56 (1991)
14. Newell, A.: The Knowledge Level. Artificial Intelligence 18, 1-20 (1982)
15. Pan, J.Y.-C., Tenenbaum, J.M. & Glicksman, J.: A framework for knowledge-based computer-integrated manufacturing. IEEE Trans. Semiconductor Manufacturing 2(2), 33-46 (1989)
16. Prerau, D.S. et al.: Maintainability techniques in developing large expert systems. IEEE Expert 5 (3), 71-79 (1990)
17. Reinder, M., Bredeweg, B.: Strategic Reasoning as a Reflective Task. In IMSA '92 International Workshop on Reflection and Meta-level architecture (1992)
18. Soloway, E. & Bachant, J. & Jensen, K.: Assessing the Maintainability of XCON-in-RIME: Coping with the problems of a VERY Large Rule-Base. AAAI-87 (1987)
19. Wachsmuth, I. & Gängler, B.: Knowledge Packets and Knowledge Packet Structures. In O. Herzog & C.-R.Rollinger (eds.): Text Understanding in LILOG: Integrating Computational Linguistics and Artificial Intelligence, Berlin: Springer (1991), pp. 380-393
20. Wachsmuth. I. & Heller, B. & Meyer-Fujara, J.: HYPERCON: Modulare Wissensbasen für Hypertonie-Konsultation. Research Proposal, University of Bielefeld, Faculty of Technology (1992)
21. Wachsmuth. I. & Meyer-Fujara, J.: Adressing the Retrieval Problem in Large Knowledge Bases (Summary). Proc. 3rd Conf. Computational Intelligence (CI-90), Milano (1990)
22. Wielinga, B. & Schreiber, A. Th. & Breuker, J.: KADS: a modeling approach to knowledge engineering. Knowledge Acquisition 4, 5-53 (1992)
23. Yourdon, E. & Constantine, L.L.: Structured Design – Fundamentals of a Discipline of Computer Program and Systems Design. Englewood Cliffs: Prentice-Hall (1979)