UvA-DARE (Digital Academic Repository)

A method for the acquisition of formalized knowledge in pathology

van Ginneken, A.M.; Jansen, W.; Smeulders, A.W.M.; van der Lei, J.; Baak, J.P.A.

Publication date
1990

Published in
Methods of Information in Medicine

[Link to publication](#)

Citation for published version (APA):

van Ginneken, A. M., Jansen, W., Smeulders, A. W. M., van der Lei, J., & Baak, J. P. A. (1990). A method for the acquisition of formalized knowledge in pathology. *Methods of Information in Medicine*, 29, 182-192.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

A078370361

NCC/IBL AANVRAAGBON

KOPIE PERIODIEK EGB

(7)

19-04-2005

Datum indienen : 18-04-2005 20:28 5493-1 Clearing House
Datum plaatsen : 18-04-2005 20:28
Aanvrager : 0004/9998
Aanvraagident :
Aanvragerident : 0004/9999
Eindgebruiker : 041631433

Telefoonnummer : 050-3635057
Cooperatiecode : R

Leverwijze : Elektronisch
Fax :
Ftp :
E-Mail : m.s.van.delden@rug.nl
Ariel :

Plaatscode : 831472189 ; CBa 1367 ; ; 1976 V15 -

(1) ☐ Origineel gestuurd	(6) ☐ Niet beschikbaar
(2) ☐ Copie gestuurd	(7) ☐ Uitgeleend
(3) ☐ Overige	(8) ☐ Wordt niet uitgeleend
(4) ☐ Nog niet aanwezig	(9) ☐ Bibliografisch onjuist
(5) ☐ Niet aanwezig	(0) ☐ Bij de binder

A078370361

NCC/IBL AANVRAAGBON

Verzamelnota volgt.
KOPIE PERIODIEK EGB

19-04-2005

Datum indienen : 18-04-2005 20:28 9560-1 UvA Keur
Datum plaatsen : 18-04-2005 20:28 UB Groningen
Aanvrager : 0004/9998 Broerstraat 4
Aanvraagident : 9700 AN Groningen
Aanvragerident : 0004/9999
Eindgebruiker : 041631433 tav

Aantal

PPN Titel : 831472189
Titel : Methods of information in medicine : international journal
: for the methodology of medical research, information and
Auteur : documentation : Internationale Zeitschrift für die
Deel/Supplem. : Methodenlehre der medizinischen Forschung, Information und
Corporatie : European Federation for Medical Informatics
Jaar/Editie : 1962 Extern nummer :
Uitgave : Stuttgart [etc.] Schattauer
Serie/Sectie :
Pag-ISSN/ISBN : 0026-1270

Plaatscode : 831472189 ; CBa 1367 ; ; 1976 V15 -

Jaar : 1990-00-00
Volume : 29
Aflevering :
Eindgebruiker : 041631433 Aanvraagident. :
Auteur : A.M. van Ginneken, W. Jansen, A.W.MUVA KEUR (UB GRONINGEN) Lei,
Artikel : A method for the acquisition of formalized knowledge in patholo
Bladzijden : 182-192
Bron :
Opmerking : arno ID: 136692

Astrid M. van Ginneken¹,
W. Jansen²,
A. W. M. Smeulders³,
J. van der Lei¹, J. P. A. Baak⁴

A Method for the Acquisition of Formalized Knowledge in Pathology

¹Department of Medical Informatics, Erasmus University Rotterdam

²Laboratory for pathology, microbiology and serology, Deventer*

³Faculty of Mathematics and Computer Science, University of Amsterdam**

⁴Department of Pathology, Free University, Amsterdam, The Netherlands

Abstract: A tool is introduced for the acquisition of pathology knowledge in a formalized form, directly by the expert. Formalization of the knowledge is intended to make descriptive pathology knowledge more suitable for computerized diagnostic support since a formal representation of knowledge allows more extensive indexing, hence more flexible access. The knowledge acquisition (KA) tool also provides a useful research instrument to investigate to what extent pathology knowledge can be made explicit, to what degree ambiguity is present, in what way experts differ when formalizing knowledge, and whether it is feasible to incrementally acquire decision criteria on the basis of the formalized descriptive knowledge.

Crucial in the design of the KA tool is the incorporated meta-knowledge, which is reflected by the knowledge-base structure and is used to elicit knowledge from the expert. Knowledge is acquired from the expert via a menu-driven user interface, which follows the general steps of the pathologist when describing a case. The paper discusses the considerations underlying the design, the implementation of the KA tool, and the research goals.

Key-Words: Knowledge Acquisition, Knowledge Formalization, Computer-Assisted Diagnosis, Pathology

1. Introduction

The knowledge acquisition (KA) tool described in this paper is developed to obtain formalized descriptive pathology knowledge directly from the expert. The need for a formal representation of pathology knowledge stems from the fact, that this knowledge in its present form, i.e., books, is diagnosis-oriented: knowledge is stored by diagnosis name. In practice the problem is often reversed: from findings to diagnosis. Therefore, a considerable searching effort may be needed to find a diagnosis which fits a particular set of findings [1]. This

means that books, though most widely used for consultation, are not the most suitable media to handle the "inverse problem" of diagnosis making.

The majority of diagnostic support systems in pathology, as well as in other fields in medicine, are developed to give diagnostic decision support on the basis of findings. Many of these systems have in common that they have a high performance in a small domain. This is due to the fact that well-defined criteria for distinguishing diagnoses are easier to formulate for a small domain, where specific tests or features permit reliable conclusions. Examples of well-known reasoning strategies, used in decision support or expert systems, are rule-based reasoning, criteria tables, decision trees and frame-based reasoning.

Rules, for example, are used in MYCIN and in a diagnostic system for

leukemia [2, 3]. The rules in these systems involve the results of very specific laboratory tests. Rule-based systems are also used in quantitative pathology to interpret results [4, 5], but as such they are mostly used to obtain additional diagnostic information: a diagnosis has to be made on the basis of microscopic inspection to justify additional tests. Many diagnostic criteria in pathology cannot be expressed in numbers and are more difficult to capture in rules. In addition, rule bases are not suited for direct consultation and they become difficult to maintain when they are large [6].

AI/RHEUM uses criteria tables to make diagnoses in the field of rheumatology [7]. Each diagnosis is made on the basis of a minimum number of major and/or minor findings, which have to be present in the patient. Knowledge of this form may

* Formerly Department of Pathology, Free University Amsterdam.

** Formerly Department of Medical Informatics, Erasmus University Rotterdam.

be available for specific problems in pathology, but it is not (yet) available on a larger scale.

The same argument holds for TEGUMENT, a system for dermatopathologic diagnosis [8]. Following a decision tree, findings are used stepwise to assign a case to a decreasing set of diagnoses until, at a leaf of the tree, a final diagnosis is made. At present, 8 or 9 steps at most are required to arrive at a diagnosis. However, the complexity of such a tree will rapidly increase at the expense of maintainability, when the domain becomes larger.

In INTERNIST and PATHFINDER, frames are used to represent disease profiles which list for each diagnosis a number of relevant features [9, 10]. Weights, assigned to each feature, are used by the reasoning system to produce a differential diagnosis by descending likelihood. In these two systems the disease profiles provide information indexed by diagnosis name as well as findings. However, since the findings are represented in PATHFINDER as a list there is no formal expression of how the tissue components are spatially related. Although the pictures convey this information, it is implicit: it is not reflected in the knowledge-base and cannot be accessed as explicit verbal diagnostic criteria.

In summary, there are several high-performance diagnostic support systems, but their applicability in the field of pathology is limited and they are not at all uniform in their knowledge representation and inference strategies. For a system to become widely used it is crucial that it covers a large portion of the pathology domain.

In most expert systems, the knowledge-base is built with the aid of a knowledge engineer, who transfers the knowledge of the expert into the structured format of a computer system. Essential to this process is that the knowledge engineer has an understanding of the domain of application; otherwise, part of the knowledge may be distorted, lost or not even elicited from the expert [11, 12]. However, skill in both computer science and the field of application is rare and takes

years to acquire. Because of the knowledge engineering bottleneck, a different strategy for knowledge acquisition has emerged: learning from well-documented cases [13]. However, the learning process in these systems usually is a long-term process, which is dependent on the available material. In addition, experts have little or no control over the completeness of the acquired knowledge. Therefore, it is desirable that the expert can fill a knowledge-base by direct interaction with the computer. The development of a tool for the acquisition of formalized knowledge requires knowledge engineering at the meta-level as there has to be information about which terms are allowed and which relations among them are legal in the domain. The terms and their relations determine which knowledge can, in principle, be expressed. The meta-knowledge provides a structured environment for knowledge acquisition, which not only helps organized knowledge entry, but also improves the consistency and completeness of the knowledge-base [14].

We have developed a tool for the acquisition of pathology knowledge in a formalized form directly from the expert. Although the meta-knowledge was restricted to pathology of the ovary, the design permits the definition of meta-knowledge and consequent acquisition of knowledge for other parts of pathology as well. The paper discusses the considerations underlying the design of the KA tool, the research objectives with the tool, and its implementation.

2. Considerations Underlying the Design of the KA Tool

When building a user interface for knowledge acquisition, the question arises whether or not to use existing software. Expert system shells share the problem that they are seldom suited for direct use by experts in the field of application. They have to be general and flexible on the one hand, and sufficiently simple to operate on

the other. The result is a reduced expression capability. Shells, which combine both flexibility and a large expression capability have a complex user interface and require programming skills of the user [15]. Shells with simple, self-explanatory user interfaces are usually rigid and suitable for applications in well-circumscribed and restricted domains [16, 17]. Furthermore, several shells are derived from an existing application, like EMYCIN from MYCIN and PROTEGE from OPAL, and are most suited for knowledge acquisition in domains closely related to the original field of application [18, 19]. As we focused on acquisition of formalized descriptive knowledge as a subject worth studying in its own right, we did not want to restrict ourselves to an available shell.

Insight into the general diagnostic work-up of a case is necessary to organize knowledge acquisition such that an expert can enter features in a familiar order. As can be seen in the clinical environment and in textbooks of pathology, the general work-up is from macroscopic structure to microscopic detail. The same order is adopted in the design of the KA tool.

The expression of pathology knowledge in formalized features requires consideration of several topics, which will be discussed briefly.

Experts do not always use the same terminology. Synonymy poses no problem, but homonymy does. Homonymy is the use of the same word for different intentions, which implies a restricted consensus about the meaning of those words. However, the use of books bears the same problem: different readers may interpret a single text in different ways.

For the purpose of indexing on features, it is important to minimize ambiguity in the knowledge-base: one pathologist should not follow a variety of ways to express a specific feature. Ideally, each feature should map on one expression only in the knowledge-base. We have chosen to minimize ambiguity and the use of homonyms by offering a restricted vocabulary.

As to consensus it is important to realize that it may vary with the level of detail in which the knowledge is expressed. In general, no more detail

should be expressed than is diagnostically significant.

The explicit expression of pathology knowledge requires that features are placed in a proper context. An example of context is the occurrence of certain cell-types related to the tissue structures in which they occur. Context also includes the expression of the spatial architecture of tissue components: cell-type A surrounds the vessels, or cell-type A is present in a diffuse pattern.

An important issue in the explicit expression of knowledge is the topic of uncertainty. Many diagnoses are not made with 100% certainty. We have to deal with uncertainties, which are part of the knowledge itself: the fact that a finding or even a combination of findings seldom justifies a diagnostic conclusion with complete certainty. The expression of uncertainties and strategies to reason with them are extensively discussed elsewhere [20, 21]. Accurate numeric expressions of uncertainty are scarce and difficult to obtain. However, as there are no more feature combinations than diagnoses, we have chosen to focus on the qualitative expression of diagnostic information with the possibility to use a limited set of words to indicate percentages of occurrence. Consensus about the meaning of these words is encouraging for future codification and expressing them with percentages in brackets may promote their uniform use [22, 23].

Since the relevance of features may vary with each differential diagnostic problem, it is important that the knowledge about a diagnosis is as complete as possible. The aim is that the diagnostic information is at least unique for each diagnosis. In addition, what seems trivial to the expert might not be trivial to the general pathologist. To encourage the expert to be as extensive as possible in the explicit expression of knowledge, in some way all possibly relevant features should be considered. In that manner attention is drawn to that part of the expert's knowledge which would not otherwise have been made explicit [24].

Figure 1 Organization of frame definitions and instantiations in an Epitool knowledge tree. A concept and its aspects constitute a frame.

3. Implementation of the Knowledge Acquisition Tool

3.1. The Knowledge-Base

To implement the knowledge-base we used the software package Epitool [25], an environment for the development of knowledge processing software. The package is especially suited for the development of frame-based applications [26, 27].

In Epitool frame definitions are organized in a hierarchy, which represents a set-subset relationship. This implies that inheritance in the hierarchy involves the slots: each frame definition has at least the same slots as its parent in the hierarchy. Each slot in a frame definition has properties which define its type of value, and whether that value is dynamic or constant. Instantiations are always at the bottom level of the hierarchy. Each slot value has to be one or a set of instantiations, which match the type as specified in the slot property. Figure 1 shows the basic structure of an Epitool tree.

In the design of the knowledge-base the term "tree" is used to denote an implementation of an Epitool hierarchy. Two trees are in use: a classification tree and a knowledge tree.

The frame definitions at the nodes of the classification tree constitute diagnosis groups and diagnoses. We have chosen to use the classification of

the World Health Organization (WHO) for ovarian tumors as the basis for the ordering of the tree [28]. The classification tree plays a secondary role and is used for two purposes: first, to provide the expert pathologist with a familiar way to select a diagnosis and, second, to link a diagnosis description with the diagnosis name. Consequently, frame definitions in the classification tree have two slots. One slot stores the name of the expert who described the diagnosis, and the other slot is used to link the diagnosis name with the actual description in the knowledge tree. Note, that a diagnosis can have more than one instantiation in order to allow several experts to enter their knowledge about a diagnosis. A part of the classification tree is given in Figure 2.

The knowledge tree is different from the classification tree: it is meant to hold diagnostic information in the form of formalized features. When knowledge is not yet entered, the tree contains meta-knowledge, i. e., about the vocabulary and structure of the diagnostic information. Each frame definition is in fact part of the meta-knowledge and corresponds to a concept in pathology and specifies how that concept is to be described. Instantiations of a frame definition are actual descriptions of its corresponding concept within a specific diagnosis. When knowledge is entered the knowledge tree grows, its contents being represented by numerous in-

Figure 2 Excerpt from the classification tree based on the WHO classification for ovarian tumors.

Figure 3 Part of the meta-knowledge with frame definitions, and predefined instantiations. The meta-knowledge as a whole defines the syntax, semantics and scope of the descriptive knowledge to be acquired.

stantiations. As the knowledge tree can hold information concerning many different diagnoses, a frame definition may have instantiations belonging to

more than one diagnosis. The information with respect to a particular diagnosis is represented by a set of instantiations, which are linked in the

proper context by means of their slot values.

The frame definitions are divided into macroscopic and microscopic

Figure 4 Diagnosis features as represented in the knowledge tree (a) and in their functional context (b).

frame definitions following common practice. Part of the meta-knowledge is shown in Figure 3. The instantiations in Figure 3 are not specific for one diagnosis; they represent predefined instantiations, each of which occurs only once. They can be used to characterize diagnoses, but are not described themselves in further detail. Hence, slots may receive predefined instantiations as their value. Examples of predefined instantiations are "slight" and "often", which serve as values for the slots "degree" and "frequency", respectively. In this way it can be expressed how frequent a feature occurs and to what degree it is present.

The knowledge tree, as presented in Figure 4a, shows the frame definitions and instantiations in a set-subset hierarchy. It does not reflect the relations between the instantiations of one diagnosis. To visualize the formalized knowledge about a single diagnosis, the information can be presented in another way: a tree where the connec-

tions represent part-of links between instantiations. This is shown in Figure 4b.

3.2. Knowledge Acquisition and Representation

The design of the knowledge tree enables controlled knowledge acquisition: as the creation of an instantiation requires the specification of its slot values, the type of the slot values determines which frame definitions are to be instantiated next. Each slot value is always worked out down the tree. When the leaves of the tree are reached, recursion takes the user back in the direction of the stem. However, the assignment of a new slot value at a higher level reverses the direction again: in practice the user moves up and down the tree many times before the knowledge entry session finishes at the top node of the tree.

In order to achieve uniformity during knowledge acquisition, five slot categories have been defined; each slot belongs to exactly one of these

five categories. Depending on which slots are defined at a frame definition, part or all of the categories are represented. Slots receive their values in a sequence determined by the following fixed order of categories:

- 1 parameter slots,
- 2 correlation and reference slots,
- 3 property slots,
- 5 part-of slots.

The *parameter slots* express frequency and degree of a feature as well its spatial occurrence (e.g., focal or dispersed). *Correlation and reference slots* are used to express correlation with or reference to some other instantiation: when the same tubules are always surrounded by necrosis, these tubules and necrosis are correlated. References are used to avoid redundancy in the knowledge tree: A cell-type, which occurs in the stroma as well as in the epithelium, need be described only at one position in the tree and can be referenced in others. In other words: One instantiation can serve as the value of several reference slots. *Property slots* characterize instantiations as a whole, e.g., "color", "size", "appearance" and "weight". *Phenomenon slots* describe which phenomena might occur associated with an instantiation such as "necrosis" in a "stroma". Finally, *part-of slots* link an instantiation with its composing parts: the contents of a cyst or the nucleus of a cell. By means of the part-of slots the path through the knowledge tree is from macroscopic to microscopic scale.

Using an example the following topics will be addressed: the user interface for knowledge acquisition, the internal representation of the system's actions during knowledge entry and the external representation of the knowledge to the user.

A part of an exemplary knowledge entry session is shown in Figure 5. The entry of knowledge is menu-driven, thereby guaranteeing the use of the restricted vocabulary. The order in which the menus appear on the screen depends on the choices of the user. Note, that menus allow for more than one choice, which is visible in the menus "Nucleus" and "Chromatin". An additional choice can be made when the previous one is com-

Figure 5 An example of menu-driven knowledge entry. Note, that each choice in a menu (b) is followed by a fixed sequence of menus to allow for the specification of parameter slots, correlations and references when necessary (a). A separate window (c) displays feedback during knowledge entry.

pleted. The expert may leave a menu by selecting "Ready".

References to instantiations are specified as follows. As soon as the entry "Reference" is selected, the user returns to the first menu, containing the major diagnostic groups of ovary pathology. As usual, the user then progresses through the classification tree with the difference that menus show only diagnoses already entered in the knowledge-base by the same author. As soon as a selection is made, the instantiations characterizing the selected diagnosis are displayed as nodes of a tree (see Appendix). One

of them can be selected with the mouse. With the selection of an instantiation from the tree, the specification of the reference is completed. The specification of a correlation with an instantiation is done in the same way, with the difference that the user also specifies the type of correlation like "resembles", "surrounded by" or "intermixed with".

The formal representation of knowledge in the knowledge-base is hidden for the user. There is an one-to-one correspondence between the knowledge entered by the user, and its formal representation. Table 1 shows

the formal representation of the actions taken in Fig. 5. Note, that instantiations are only made when the frame definition has no subtree of frame definitions in the hierarchy.

The form of the feedback to the user depends on the situation. Feedback during knowledge acquisition is straightforward and consists of the chronological display of the choices made by the expert. These choices, however, are displayed with varying margination to clearly visualize their context. Both the feedback and the menu sequence during knowledge acquisition are visible in Fig. 5. The feedback is necessary to provide the user with sufficient overview; at the present stage of development only one menu is visible at a time.

For the purpose of inspection of the knowledge-base, the information can be retrieved in the form of a tree as shown in the Appendix. Note, that this is the same kind of tree which is displayed when specifying a correlation or a reference.

Table 1 Sequence, showing the relation between the actions of the user and the formal representation of the operations on the knowledge-base during knowledge entry as shown in Fig. 5. A hyphen ("—") indicates that no operation in the knowledge-base is performed. The text in boldface denotes actions performed on the knowledge-base.

Menu selection	Operation in knowledge-base
epithelial layer	Create Instantiation of epithelial layer: obj10
ready	—
non-standard cells	Create Instantiation of non-standard cells: obj11 Set Value obj10.non-standard cells = obj11
ready	—
form	— No instantiation is created of form, as it has a subtree of frame def's.
polygonal	Create Instantiation of polygonal: obj12 Set Value obj11.form = {obj12} The slot form may have several values, hence the values are stored in a set.
ready	—
nucleus	Create Instantiation of nucleus: obj13 Set Value obj11.nucleus = obj13
ready	—
chromatin	— No instantiation is created of chromatin as it has a subtree of frame def's.
hyperchromatic	Create Instantiation hyperchromatic: obj14 Set Value obj13.chromatin = {obj14} The set represents that the slot chromatin may have several values.
degree	— No instantiation is created of degree, only of its predefined instantiations.
little/slightly	Set Value obj14.degree = little/slightly
ready	—
normochromatic	Create Instantiation normochromatic: obj15 Set Value obj13.chromatin = {obj14, obj15}
ready	—
pattern	— No instantiation is created of chromatin, as it has a subtree of frame def's.
pleiomorphic	Create Instantiation pleiomorphic: obj16 Set Value obj13.pattern = obj16
ready	—
ready	—
ready	—
ready	—

4. Research Objectives

Prior to building a large knowledge-base, which can be used for diagnostic support, it is important to assess to what degree pathology knowledge can be made explicit in our knowledge-base and to what degree the knowledge-base structure is ambiguous.

It is generally understood that pathology knowledge is not completely expressible in words only; part of the knowledge is conveyed in pictures. Hence, the question can be narrowed down to what degree the verbal knowledge can be made explicit. The knowledge, as laid down in textbooks, is mainly descriptive, whereas experts know in addition how to use the descriptive knowledge for solving diagnostic problems. As explained earlier, the latter type of knowledge, in the form of structured decision criteria, is at present only available in small parts of pathology. For the purpose of formalizing pathology knowledge on a large scale, focus is on the knowledge in the literature and the expert's descriptive concept of diagnoses.

In the knowledge-base it is undesirable that a single feature can be expressed in more than one way as this makes indexing of features much more complex and prone to inconsistencies. Ambiguity in the knowledge-base structure can be detected by comparing descriptions of the same diagnosis, entered by different pathologists. The degree to which formalizations of one text by different people are similar is a measure of the unambiguousness of the mapping from the text onto its formal representation. However, it is not expected that ambiguity can be avoided completely. The presence of separate concepts for specific features, which can also be formalized using a combination of other concepts, is sometimes desirable since the level of detail of the formalizations, which would be required otherwise, entails an even greater ambiguity.

Though the formalization of pathology knowledge would be a major step forward to opening up pathology knowledge for diagnostic support, the next step in research would be to investigate the differences between the knowledge of the expert and the knowledge as laid down in textbooks. Asking experts to formalize their knowledge by heart with the KA tool may reveal what kinds of knowledge the expert has in addition to books. When part or all of their formalizations show more information than those based on books, there is an indication that experts, at least sometimes, use more descriptive information in diagnosis making than is expressed in books. Comparing the formalizations of more than one expert may reveal to what degree such additional features are considered important by fellow experts. There are studies which reveal that an important difference between expert and novice is the fact that the expert needs less information and less processing to support his or her decision [29]. In other words, the expert can select from a large amount of information those features which are relevant for a conclusion. The expert knows how the pieces of information are related and what part they play in the decision process. A study in the field of congenital heart diseases showed that

consensus among experts was highest with respect to those features which were regarded most important [30].

Following from the previous, the next step is to elicit from experts what role the various formalized features play in diagnostic decision making. Using the graphical trees made by the KA tool, experts can indicate for each feature its relevance for decision support, such as: essential for the diagnosis, sufficient to confirm the diagnosis, or sufficient to reject the diagnosis. In this way diagnostic support can be developed in two phases: formalization of descriptive pathology knowledge and stepwise acquisition of knowledge for decision support.

5. Current Status of Development

To find out whether or not the KA tool is sufficiently developed to do the intended research, a pilot study has been conducted with three experts in the area of ovarian pathology. The experts were asked to characterize a dysgerminoma by heart and a Sertoli cell tumor from an existing text. Since none of the experts was familiar with the system, the sessions were preceded by an introduction of approximately 20 min. During the experiment the experts were free to ask questions about the use of the system or to comment on it. The Appendix shows two trees, representing the knowledge about a dysgerminoma, entered by two different experts. The experiment was intended to gain insight into the user-friendliness of the system, the expression capability of the KA tool, and the mental process of formalizing the knowledge.

As to the operation of the system, none of the experts had difficulties with understanding how to enter knowledge. Initially, the participants had to try some menus to find out whether or not these contained the term they needed. This is explained by the fact that some terms can be interpreted in more than one way: "smooth" can be viewed both as an appearance and as a touch. Only one

interpretation is used in the system. At present, one menu is visible at a time, but the participants preferred the display of a few previous menus as well for a better overview. The average time needed to enter a diagnosis was 40 min, which is expected to shorten as users become more familiar with the system.

The degree to which the knowledge could be made explicit was considered sufficient. None of the experts had the feeling that features were left unspecified because they were too complex or too subtle to express in the syntax of the system. Only at two instances a choice was felt missing in a menu. Additional options are expected to be missing when other diagnoses are entered, but these gaps can be easily filled by the addition of new, non-overlapping choices.

The order, in which issues passed in revue, was experienced as natural during knowledge entry by heart. In contrast, when entering formal knowledge on the basis of the existing text, they experienced the conversion of the ordering of the written knowledge to the ordering of the system as a separate mental effort. Several times during one session the pathologist searched the entire text to find features requiring a selection in the menu on display. Though a study is required with many more diagnoses, it was striking that the formalizations made from the text were identical, whereas the formalizations by heart showed some differences as can be seen for two trees in the Appendix.

This small experiment allows for no conclusion other than that the present state of development justifies research on a larger scale.

6. Conclusion

A method for the acquisition of formalized pathology knowledge directly by the expert has been described. Crucial in the method is the use of knowledge at the meta-level to guide the expert in the process of knowledge formalization. The meta-knowledge is incorporated in the design of a KA tool in such a way that it

is reflected by the structure of the knowledge representation and is used for a menu-driven user interface. This approach guarantees that the pathologist automatically uses the proper syntax and vocabulary. As the process of formalization is guided conform the usual work-up in pathology, completeness is promoted.

The advantage of knowledge engineering on the meta-level was twofold. First, implementation and testing of the meta-knowledge required only a few days, which is far less than the time needed to acquire the actual descriptions of numerous diagnoses. Second, the skill required from the knowledge engineer shifted from detailed knowledge of the domain of application to more generally applicable knowledge of how to elicit from the expert the structure of the knowledge, and the general steps in the diagnostic work-up.

Though the naked knowledge tree will have to be extended with many more concepts to allow for the formalization of pathology on a larger scale, formalized pathology knowledge is a *conditio sine qua non* for access to knowledge, based on findings. There will be no problem of scaling up to the entire field of pathology as each organ can be treated separately: It is always known from where patient tissue is taken and histological abnormalities are usually spatially separated, each posing its own diagnostic problem.

The KA tool provides a useful tool for further research on the feasibility of formalization, the differences among experts in formalizing knowledge, and the feasibility of incremental acquisition of decision criteria on the basis of existing formal descriptive knowledge.

Acknowledgements

This work has been supported by grant PF 28-1207 of the Praeventiefonds. We gratefully acknowledge the participation of C. Kooijman M. D., Ph. D. and J. A. J. Spaas M. D. in the pilot study.

Appendix

Two trees, each representing formalized knowledge about the diagnosis dysgerminoma in its functional context. When specifying a reference to an instantiation the user can select a node from this tree. Nodes representing values of parameter slots cannot be referenced.

The two trees reflect several kinds of differences in the formalized knowledge. Many differences in extensiveness reflect differences between the

Figure 6 Tree, containing formalized knowledge about dysgerminoma.

Figure 7 Tree, containing knowledge about dysgerminoma, also containing features that should not be present in the diagnosis.

pathologists as to what features should be mentioned. For example, the second tree mentions more tissue structures than the first. Furthermore, the first tree contains features, which (may) fit with a dysgerminoma, whereas the second tree also contains features, which should not be present in the diagnosis. Differences in the characterization of corresponding features are found in the form of the nuclei and several values of parameter slots. Finally, the two trees show the presence of ambiguity in the naked knowledge tree: the presence of inflammation can be characterized by the presence of an inflammatory infiltrate, by mentions of lymphocytes in the tissue structures or by a combination of both.

REFERENCES

- [1] van Ginneken AM, Baak JPA, Jansen WJ, Smeulders AWM. Evaluation of a diagnostic encyclopedia workstation for ovarian pathology. Submitted for publication in Hum Pathol.
- [2] Buchanan BG, Shortliffe EH. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading MA: Addison-Wesley, 1984.
- [3] Alvey PL, Graves ME. Observations on the development of a high performance system for leukemia diagnosis. In: Bramer MA ed. *Research and Development in Expert Systems III*, Cambridge University Press, 1986: 99-110.
- [4] Conference digest of: The First International Conference on Artificial Intelligence Systems (EXPERT SYSTEMS) as Diagnostic Consultants for the Cytologic and Histologic Diagnosis of Cancer. Universal City CA, 1987.
- [5] Conference digest of: The Second International Conference of Artificial Intelligence Systems (EXPERT SYSTEMS) as Diagnostic Consultants for the Cytologic and Histologic Diagnosis of Cancer. Chicago ILL, 1988.
- [6] Hayes-Roth F. Rule-based systems. Comm ACM 1985; 28: 921-32.
- [7] Kingsland LC, Lindberg DAB, Sharp GC. Anatomy of a knowledge-based consultant system: AI/RHEUM. M.D. Computing 1986; 3: 18-26.
- [8] Potter B and Ronan SG. Computerized dermatopathologic diagnosis. J Am Acad Dermatol 1987; 17: 119-31.
- [9] Miller RA, Pople HE, Myers JD. INTERNIST-I, An experimental computer-based diagnostic consultant for general internal medicine. N Eng J Med 1982; 307: 468-76.
- [10] Horvitz EJ, Heckerman DE, Nathwani BN, Fagan LM. Diagnostic strategies in the hypothesis-directed PATHFINDER system. In: *Proceedings of The First Conference on Artificial Intelligence Applications*. New York: IEEE Computer Society, 1984: 630-6.
- [11] Musen MA, van der Lei J. Knowledge engineering for clinical consultation programs: modeling the application area. Meth Inform Med 1989; 28: 28-35.
- [12] Gaines BR. An overview of knowledge-acquisition and transfer. Int J Man-Machine Studies 1987; 26: 453-72.
- [13] Hunter L, Silbert JA. Progress Report on IVY: A learning system for intelligent information retrieval in pathology. *Artificial Intelligence in Medicine Workshop*. Yale University New Haven CT, 1987.
- [14] Walton JD, Musen MA, Combs DM, Lane CD, Shortliffe EH, Fagan LM. Graphical access to medical expert systems: III. Design of a knowledge acquisition environment. Meth Inform Med 1987; 26: 78-88.
- [15] Gevarter B. The nature and evaluation of commercial expert system building tools. Computer May 1987; 24-41.
- [16] Vital J. Languages for use in artificial intelligence research. *Analyt Quant Cytol Histol* (accepted for publication).
- [17] Salzman GC, Krall RB, Marinuzzi JG. Knowledge engineering software. A demonstration of a high-end tool. *Analyt Quant Cytol Histol* 1988; 10: 219-224.
- [18] van Melle W. A domain-independent production-rule system for consultation programs. In: *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*. Tokyo, 1979: 260-2.
- [19] Musen MA. Generation of visual languages for development of knowledge-based systems. In: Korfhage RR, Jungert E, Ichikawa, eds. *Visual Languages and Applications*. Vol II, New York: Plenum, in print.
- [20] van Ginneken AM, Smeulders AWM. Reasoning in uncertainties: an analysis of five strategies and their suitability in pathology. *Analyt Quant Cytol Histol*, (accepted for publication).
- [21] Gammerman A, Creaney N. Modelling of uncertainty in expert systems. *Second International Expert Systems Conference*. Learned information 1986: 265-74.
- [22] Kong A, Barnett GO, Mosteller F. How medical professionals evaluate expressions of probability. N Eng J Med 1986; 315: 740-4.
- [23] Toogood JH. What do we mean by "usually"? Lancet 1980; 1: 1094.
- [24] Musen MA. *Automated Generation of Model-Based Knowledge-Acquisition Tools*. London: Pitman, San Mateo CA: Morgan Kaufmann, 1989.
- [25] Epitool Development Environment, (Copyright 1986 by Epitec AB), S:t Lars-gatan 12, S-582 24 Linköping, Sweden.
- [26] Fikes R, Kehler T. The role of frame-based representation in reasoning. Comm ACM 1985; 28: 904-20.
- [27] Minsky M. A framework for representing knowledge. In: Winston PM ed. *The Psychology of computer vision*. New York: McGraw-Hill, 1957: 211-77.
- [28] Serov SF, Scully RE, Sobin LH. Histological typing of ovarian tumors. In: *International Histological Classification of Tumours*. Vol 9. Geneva: World Health Organization, 1973: 37-54.
- [29] Kolodner JL. Towards an understanding of the role of experience in the evolution from novice to expert. Int J Man-Machine Stud 1983; 19: 497-518.
- [30] Leao BF, Rocha AF. Proposed methodology for knowledge acquisition: A study on congenital heart disease diagnosis. Meth Inform Med 1990; 29: 30-40.

Address of the authors:
A. M. van Ginneken M. D., Ph. D.,
Department of Medical Informatics,
Erasmus University,
P.O. BOX 1738,
3000 DR Rotterdam,
The Netherlands