UvA-DARE (Digital Academic Repository)

Design of a diagnostic encyclopaedia using AIDA

van Ginneken, A.M.; Smeulders, A.W.M.; Jansen, W.

DOI

[10.1016/0169-2607\(87\)90093-9](https://doi.org/10.1016/0169-2607(87)90093-9)

Publication date

1987

Published in

Computer Methods and Programs in Biomedicine

[Link to publication](#)

Citation for published version (APA):

van Ginneken, A. M., Smeulders, A. W. M., & Jansen, W. (1987). Design of a diagnostic encyclopaedia using AIDA. *Computer Methods and Programs in Biomedicine*, 25, 339-347. [https://doi.org/10.1016/0169-2607\(87\)90093-9](https://doi.org/10.1016/0169-2607(87)90093-9)

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.

CPB 00892

Design of a diagnostic encyclopaedia using AIDA

Astrid M. van Ginneken^{1*}, Arnold W.M. Smeulders^{1*} and Wicher Jansen²

¹ Department of Medical Informatics, Erasmus University, Rotterdam and ² Laboratory of Pathology, Free University, Amsterdam, the Netherlands

Diagnostic Encyclopaedia Workstation (DEW) is the name of a digital encyclopaedia constructed to contain reference knowledge with respect to the pathology of the ovary. Comparing DEW with the common sources of reference knowledge (i.e. books) leads to the following advantages of DEW: it contains more verbal knowledge, pictures and case histories, and it offers information adjusted to the needs of the user. Based on an analysis of the structure of this reference knowledge we have chosen AIDA to develop a relational database and we use a video-disc player to contain the pictorial part of the database. The system consists of a database input version and a read-only run version. The design of the database input version is discussed. Reference knowledge for ovary pathology requires 1–3 Mbytes of memory. At present 15% of this amount is available. The design of the run version is based on an analysis of which information must necessarily be specified to the system by the user to access a desired item of information. Finally, the use of AIDA in constructing DEW is evaluated.

Diagnostic encyclopaedia; Pathology; Fourth-generation software

1. Introduction

With the advent of the personal computer (PC) the last barrier may have been taken to introduce the computer at the physician's desk. There, the PC may acquire a foremost position in offering three sorts of information: information about the actual state of the patient (e.g. laboratory tests), archival information about the patient, and reference knowledge with respect to diseases. We do not discuss the possibilities of PCs in acquiring actual patient information (see, for example, [1]) nor exploit its possibilities for maintaining the patient archive (e.g. [2]). In this paper we focus on the construction of an encyclopaedia, designed to contain reference knowledge. In general, reference knowledge reflects the state-of-the-art in medicine,

indicating what is known and to be known about the disease when making decisions.

For the encyclopaedia described here we have chosen the field of pathology as in pathology everything serves to make decisions. These decisions can be evaluated, as they are based on observations, and they can be repeated using the same materials. The name Diagnostic Encyclopaedia Workstation has been given to the system, with the acronym DEW. Since pathology is a wide-ranging domain, we have at present restricted the encyclopaedia to the pathology of the ovary. The pathology of the ovary is quite well circumscribed and, as compared to most organs, complex enough to be a good representative of general pathology; the ovary contains various types of tissue and its hormonal status changes both monthly and in a lifetime.

As pathology is a profession based on visual (mostly microscopical) observations, it was felt absolutely necessary to include pictures in the encyclopaedia. The encyclopaedia has been real-

Correspondence: AIDA Group, Department of Medical Informatics, Erasmus University, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands.

* Formerly at the Free University, Amsterdam.

ized in the form of a database, thus containing pictorial, verbal and numerical information. In the sequel of the paper, the database design and user interface are described. The fourth-generation software system AIDA [3] was used for database design.

2. System specifications

In order to develop technical specifications for the database structure, the user interface and the required hardware, let us analyse the information needed in the diagnostic decision process in pathology.

Once patient material has been sliced and chemically treated to visualize tissue components, the microscopical slice is given to the pathologist. In combination with previous information about the case, visual inspection of the slide then leads to the diagnosis by comparing the observations with reference knowledge.

The reference knowledge consists of all the general information of pathology available to the pathologist. Reference knowledge comprises the training received, the experience gained throughout the years in similar cases, handbooks and atlases of pathology specialties frequently consulted during practice, and the expertise of consulted colleagues. Among these four components of reference knowledge a digital diagnostic encyclopaedia offers the equivalent of handbooks and atlases, with a potential expansion to training facilities and consultation of colleagues using telecommunication.

For a specific case, the reference knowledge obtained from books has the following aspects:

- (1) It is essential to have a description of the microscopic image of a distinct diagnosis, placed within the context of general characteristics of the disease. Pictures, mostly microphotographs of selected cases, provide pictorial information on the pathology of the disease.
- (2) Apart from the factual description of a diagnosis, potential confusions with other diagnoses (known as differential diagnoses) are also indicated. This helps the reader in deciding among confusing possibilities by the specification of dis-

criminating features among differential diagnoses. Simultaneously, the reader's attention is directed towards the existence of these alternative diagnoses.

- (3) Not so often found in pathology handbooks these days are illustrated case histories. However, in conveying information beyond the written, these case studies offer the reader a lively and instructive view on a particular case.

- (4) To confirm a diagnosis or to discriminate among differential diagnoses, a list of available laboratory techniques is helpful. Additional techniques in pathology encompass a cytological analysis of loose cell samples, electron microscopy of sub-cellular details, alternative chemical treatments visualizing different tissue components, quantification of microscopical features and immunopathology responses.

As mentioned above, these main components of reference knowledge can be found in the basic literature. There are, however, a number of restrictions to the use of books.

- (1) Pathology books contain a few hundred pictures at most—generally one for each diagnosis—but usually more pictures are desired to show alternative laboratory techniques at several optical magnifications. Owing to economical restrictions in publishing, especially where colour pictures are concerned, illustrations per diagnosis are limited to a few.

- (2) A book permits only a one-dimensional ordering of material; reading from the beginning to the end. Contents and indexes provide the possibility of jumping to a certain place in the book, but from thereon the order of the book must be followed. Combination of information available at different places in the book can only be achieved by boring text repetitions in the book or by specifying irritating internal references to the reader. Thus, in a book there is no satisfactory way to bring order in the information along more than one line.

- (3) The range of the field of pathology is so wide that, in order to keep a volume practicable to the reader, information has to be restricted.

- (4) The supply of information in a book cannot be adjusted according to the need of the user. This need may change with the experience of the user

in the specific field and also depend on the situation. As to the situation, the user may need supportive arguing to a diagnosis, interrogative arguing to help find one, or checklists of features in differential diagnoses.

These four restrictions play no part when a (personal) computer system is used to structure and present the information. The inventory of these demands has been translated into computer terms as follows:

- (1) A large amount of pictorial data must be stored, preferably in colour. Storing a picture digitally would require storage of some 512×512 pixels in 3 colours each, yielding 0.75 Mpixels, per image. Obviously at the present state of the art, a thousand pictures or more cannot easily be contained on a manageable computer disc. Therefore the pictorial information is stored on an outside medium with large memory capacity, whereas the verbal and numerical information are stored on disc.
- (2) Apart from pictorial information the database must contain verbal and numerical information. There has to be an efficient way to store text of variable size, and multiple indexing must be allowed for easy access of information.
- (3) The database must be large as the amount of information to be stored encompasses many books.
- (4) The software must be flexible enough to allow for database retrieval adjusted to the competence of the user.
- (5) In addition, the system must run on a commonly used personal computer of moderate cost in order to compete with the total price of books and journals.

On the basis of these demands, we have selected for the present development the IBM AT-compatible personal computer, interfaced to a video long disc player (VLP). A VLP typically costs \$1000 at present and holds a replaceable disc. The disc is randomly accessible and contains a maximum of 54,000 colour video images. Each image on the disc can be addressed by a computer using a standard RS-232 interface. Both the computer and the VLP have their own monitor. The quality of reproduction as a still picture is comparable with the best illustrations in regular books. The IBM AT at present is equipped with 600 Kb memory

and a 20 Mb Winchester drive, but it is open whether so much memory capacity is really needed for an operational system.

For database functions for an encyclopaedia, it should be emphasized that it will be used for read-only access, once constructed. Therefore, the development of the encyclopaedia can be divided into two phases. The first is the construction of a database and a program to insert the data, resulting in what we will call the input version. The second phase comprises the development of the retrieval part, which we will refer to as the run version. Adding data to the database is in principle done only once and may be somewhat laborious, but it is worth the effort if great attention is paid to the accessibility of the data in the run version.

3. Design of the database

3.1. Input version

Prior to designing a database, the optimal structure of the database must be determined. The structure of reference knowledge is basically hierarchical: the pathology of an organ encompasses some major diagnostic groups, each of which comprises several diagnoses and in a number of cases these diagnoses have several gradings. In general, we treat gradings of a diagnosis as separate diagnoses. To every diagnosis there are several case histories, literature references and a list of differential diagnoses.

Every illustration is an element of a case, but some of these also serve to illustrate more general statements or diagnoses. It may even be so that entire case studies serve to illustrate more than one diagnosis. This also holds for literature references and implies that items may have multiple connections with higher levels. However, in a hierarchical tree such connections are not allowed. Viewed upon strictly hierarchically such items are present more than once, as shown in Fig. 1. Thus, if a strictly hierarchical database structure is selected, redundancy is inevitable in the database, which is impractical in constructing the input version. The presence of redundancy can be pre-

Fig. 1. Structure of the information taken as the starting point for the construction of the database. Note that some of the items appear at more than one place in the database; in this figure, these places are marked by an 'x'.

vented by splitting the tree into a number of sub-trees and constructing a file system of reference tables. However, a hierarchical database management system does not deal with the maintenance of such file systems. Another point to mention is that a hierarchy is not the only desirable way of access to the information. Particularly when dealing with differential diagnoses, one wishes to have quick access to other diagnoses, irrespective of the location they have in the hierarchy.

The disadvantages mentioned above play no part in a relational database structure, in which all information, irrespective of its meaning, is stored in the form of relations and maintained by the database management system. This leads us to prefer a relational database over a hierarchical one. Given this choice, it must be possible to define one-to-many relations, i.e. relations must allow for composite keys. The database system must also support the feature of free record length and free field length owing to the presence of many items containing plain text of unpredictable dimension. Moreover, in favour of retrieval speed, the database system must offer the possibility of defining index relations.

Another point to consider in the selection of a suitable database system is the foreseen size of the database. A rough estimate of the size of the database can be made on the basis of its most dominating factors: the number of diagnoses, the

number of cases and the average amount of information per diagnosis and per case.

According to the WHO there are 8 major diagnostic groups in the pathology of ovarian tumors [4]. Some of these major groups can be subdivided, making a total of 25 diagnostic groups encompassing approximately 90 diagnoses, including rarities. Blaustein [5] recognizes 4 diagnostic groups in non-tumor pathology of the ovary, covering 62 diagnoses, including many rarities. Among pathology books there is some variation in these numbers, depending on the way the information is classified and ordered. A basic encyclopaedia, leaving rarities out, will thus contain 80 tumor diagnoses and 35 non-tumor diagnoses, including a description of the normal ovary. The average amount of text might be 3 pages per diagnosis and one page per case history. Taking 3 case histories as the average per diagnosis and estimating the amount of memory needed to store a page of text at 2.5 Kbytes, 115 diagnoses and 445 case histories will require approximately 2.5 Mbytes of memory for ovary pathology. It should be added that not all information has the same priority at a particular moment and need not, therefore, all be kept in central memory.

The fourth-generation software (4GS) system AIDA [3] meets all requirements for the construction of the input version. Especially the possibility of inserting free text in a relational database system and the availability as a toolkit has greatly influenced the choice. A full discussion of the choice against other candidates is presented elsewhere [6]. AIDA is a 4GS system written with MUMPS as host language. The virtues of MUMPS as an interpreter in text string manipulation and sparse matrix storage have been exploited in the development of AIDA. The AIDA-MUMPS combination therefore is ideally suited for database development.

Having specified the requirements for the database, we will now give a brief description of the actual implementation of the input version using AIDA.

For the input version the database contains a total of 14 relations; 34 numerical-type items are defined of which 5 are sequence numbers needed as part of composite keys in one-to-many rela-

tions. There are 41 character-type items of which 17 represent categories with an expected average length of 40 characters. The remaining 24 character-type items are free text. A considerable part of the free texts will remain empty as the text is not relevant or not available. It should be mentioned that in MUMPS empty items require no storage space. The input program comprises 34 different input screens for adding and correcting data. The screens appear in hierarchical order. It is possible for the user to bypass screens, which is desirable if the information being asked for is absent or irrelevant to a particular diagnosis. Basic validation per item is standard in AIDA, so there was no need to program this. Validation by comparing information among items can only be done by a computer program in a minor portion of the items and has not (yet) been implemented. Fig. 2 shows the database structure.

entered in the database and approximately 1000 pictures are available on a video disc.

The above-specified database was generated in three months (from the start, after a MUMPS course of one week, until it was released for input) by one programmer with some support of the AIDA development team. The code at its present state required 450 lines of MUMPS code apart from the AIDA toolkit.

3.2. The runtime version

An encyclopedia system is read-only once put together. As has been pointed out above, this enables separate treatment of the design of the user interface for runtime use.

In an encyclopaedia system containing a great deal of information, the presentation of data is of more than usual importance. The design of the user interface is based on an analysis of the items of information which necessarily need to be specified by the user for the system to respond and how this information can be specified with a minimum number of keystrokes.

A user who wishes to access a particular item of information about a diagnosis in a diagnostic encyclopaedia must provide six items of information. He must first specify the diagnosis to the system. As organs, diagnostic groups and diagnoses cannot be displayed on a screen simultaneously the specification of a diagnosis requires three sequential interactions: the specification of an organ (at present only one), a diagnostic group and a diagnosis. At the level of a diagnosis, the amount of available information in the encyclopaedia is too large to be displayed on a screen, so the information has to be offered partially as optional items, some of which can be displayed by default. The specification of one of these optional, non-standard items requires a fourth, different type of interaction. Two more interaction types must be offered to the user: the return of the system to its starting position and the specification of the default settings of the system. The latter covers the level of the information which is displayed, and also the choice of the default item which the system will show at the diagnosis level.

For specifying these choices to the system, there are two possibilities: user-entry by typing or selection from a presented list. For the organs a presented list requires the space of one or possibly two screens. For the time being we have restricted ourselves to the pathology of the ovary, but for pathology in general the number of items in the list depends on the degree to which the organs are subdivided. For example, in the digestive tract many parts can be distinguished, such as esophagus, stomach, liver and intestine. These parts can be treated as separate organs, thus increasing the total number of organs. As nomenclature in pathology comprises many synonyms and different spellings the obvious choice for organs is to present names to select from, rather than asking for these to be typed. When displaying it is clear which organs the system supports and how they are classified. It should be mentioned that for each of the organs only one name has to be kept in computer memory. The same argument holds for interaction two and three: the diagnostic groups and the diagnoses. Option specification could be performed by typing one or two characterizing letters, selecting from a presented list, or both. Options should preferably be offered on the same position of the screen and only when the corresponding information is available. The return of the system to its starting position, a basically different interaction, can best be performed by a single keystroke, preferably on a fixed and always visible place on the screen. Finally, specification of the default settings can best be performed by selection from a presented list of possibilities, offered on request when starting interaction with the system.

Based on the above analysis, we will now describe the implementation of the run version. As only one organ is covered at present, the system starts with the presentation of the diagnostic groups in a numbered list. A selection can be made by entering a number from the list. The system then shows a similar list, containing the diagnoses belonging to the selected diagnostic group. A diagnosis is specified in the same way as the diagnostic group. At the diagnostic level the options can be divided into two groups: general options to switch to another item of information

and specific options to obtain more detailed information about the current item. Option names hold throughout the system, but may not all be offered, depending on their relevance in a particular situation. An option such as electron microscopy is only offered when microscopic description is the current item and will not be present when the item 'prognosis' is displayed.

As to the layout of the screen, information about a selected item is displayed at the right-hand side of the screen, whereas the options are shown on the left. These screen locations for information and options are the same for all screens. The options are displayed as labels with an answer field. The labels appear in reversed video. One can pass through the fields with the arrow-keys. Options can be specified by entering <CR> or typing one or two characters in the selected answer field. Fig. 3 shows the layout of the first screen. The Figs. 4 and 5 show the default and an optional item, respectively, at the diagnosis level.

The screens as specified above have been generated with AIDA and give a good impression of the possibilities offered by the AIDA screen generator FRAMAID. Although a user interface can be built easily and quickly with the aid of FRAMAID, screen design is basically limited to the configuration of labels and their answer fields. Graphics and local scrolling options are not available in AIDA itself.

For an information reference system, the overall response time is very important. The tolerated elapse time is set to the maximum considered to be still convenient in interaction, i.e. one second. The total time needed by the system to come up

```

      Organ and its sub-classes
      OVARY
      1  1 SEROUS TUMOURS
          2 MUCINOUS TUMOURS
          3 ENDOMETRIOID TUMOURS
          4 CLEAR CELL TUMOURS
          5 BRENNER TUMOURS
          6 GERM CELL TUMOURS
          7 TUMOUR LIKE CONDITIONS

      — Select numbers
      — F (orward scroll)
      — B (ackward scroll)
      — S (top)
  
```

Fig. 3. First screen of the run version.

Microscopy: common histology	
Electron microsc.	Diagnosis: SEROUS SURFACE PAPILLOMA
Immune pathology	
Quantitative path.	
Ess. clin. data	1 An ovarian, serous, benign tumour consisting of
Ext. clin. data	papillary projections on the ovarian surface.
Macroscopic descr.	
Cytology	The surface contains papillae [6], which are almost
Diagn. criteria	completely composed of stroma [7], sometimes with
Clinical quest.	a collagenous and/or markedly oedematous [8]
Differential diagn	appearance.
Case history	
Image (nr.)	The epithelium/stroma ratio is 0.25-2.0 (vol. %
Literature (nr.)	epith. 20-70).
F (orward)	The intrapapillary spaces may contain mucin.
B (ackward)	
R (start)	The cubic epithelium of the papilloma is like that
S (top)	of ovarian surface epithelium [10]. Focally cylin-

Fig. 4. Example of a screen of the run version at the diagnosis level. The microscopic description is displayed by default. Images are displayed upon selection of numbers in the text.

with a picture is the sum of the response time of the VLP and the retrieval time of the database. As to the VLP (a Philips VLP 835) the command 'next picture', which will often be used, takes 40 ms to be carried out, corresponding to the normal speed of 25 images per second. The response time of our VLP to come up with an image at a position n , relative to the current location, is specified in Fig. 6. The times given are an average as the response is not always the same for the same jump. Remarkably enough, the figure shows that the response time is neither constant, nor does it increase linearly with n . Note also the considerable time required to come up with pictures in the vicinity. The maximum response time of the player is 4.3 s, i.e. when jumping from position 1 to position 54000. Within one organ, to which the player head may be set as soon as it has been specified, n may reach a maximum of several thousands. When a jump reaches 1000 relative

Essential clinical data	
Ext. clinical data	Diagnosis: SEROUS SURFACE PAPILLOMA
Demography	
Chemistry	
Röntgenology	
Staging	1 Usually an incidental find during operation.
Classification tree	
Therapy	
Prognosis	
Macro description	
Micro description	
Diagn. criteria	
Clinical quest.	
Differential diagn	
Casehistory	
Literature (nr.)	
F (orward)	
B (ackward)	
R (start)	
S (top)	

Fig. 5. Screen displaying the essential clinical data. Note that, compared to Fig. 4, some of the options have altered.

Fig. 6. Response time of the video disc player, needed to come up with a picture n positions beyond the current position.

positions, the response time of the player exceeds 1 s. However, the majority of jumps will be within one case history and, to a lesser degree, within one diagnosis. Only when a different diagnosis is chosen, especially when this diagnosis belongs to a different diagnostic group, may a large jump occur. Taking this into consideration, it can be concluded that the access time of the player may sometimes form a limitation in attaining the requested response time of one second. Obviously the performance of the player as described above reflects the present state of commercially available equipment.

As to the retrieval program, response time can be shortened by keeping a part of the database in central memory. In optimally defining the contents of this part, the system must anticipate which information a user is likely to ask for next. Considering the interactions needed to access information about a diagnosis, as pointed out above, the choice is obvious: the organs and diagnostic groups reside in central memory permanently. As soon as a diagnosis is selected, all information about this diagnosis, its differential diagnoses and their case histories can be kept in central memory as well. The amount of information to be kept in central memory can be estimated. Taking 5 differential diagnoses as the average, 6 diagnoses and 18 cases have to be stored in central memory, making a total of 36 pages of text, requiring approximately $36 \times 2.5 \text{ Kbytes} = 90 \text{ Kbytes}$. These memory requirements can presently easily be met and do not form a limitation in the optimization of the response time.

4. Discussion

An analysis of the reference knowledge to be covered by an encyclopaedia for pathology reveals a basically hierarchical structure, but at several levels in this hierarchy items have multiple connections with higher levels. In a hierarchical database this implies the presence of redundancy. As this is impractical in the construction phase of the database a relational database structure is chosen to develop the encyclopaedia. Among relational database software packages AIDA has been chosen, owing to its availability as a toolkit and its flexibility in handling texts of variable size and in redefining the database structure. The AIDA toolkit offers the possibility of defining a relational database with relatively little effort. The input version was constructed in three months after an initial course in the MUMPS programming language of one week. A user interface can be built easily and quickly owing to the presence of a screen generator and screen handler, but the possibilities for screen design are limited to the configuration of labels and answer fields. Graphics and scrolling in multiple windows are only available in MUMPS by the use of external function calls.

The design of the user interface for the run version is based on a fundamental analysis of the most efficient way for the user to pass on information to the system. Where the specification of an organ, diagnostic group or diagnosis is concerned, the user is offered a list of names, from which one item has to be selected. The presentation of these names in a list avoids the problem of dealing with synonyms and different spellings, and informs the user about the contents of the encyclopaedia. As the amount of information per diagnosis is too large to fit on a screen, the information is partially offered as optional items. These items can be accessed by selection of a corresponding field, typing a characterizing letter or both.

At present the encyclopaedia is restricted to pathology of the ovary. 6 diagnostic groups, encompassing 22 diagnoses and 50 cases have been entered in the database and on a videodisc. In the coming year we plan to have a basic encyclopaedia without rarities, containing 80 tumor diag-

noses, 35 non-tumor diagnoses and a description of the normal ovary.

For the run version, a response time of one second cannot be reached in all cases, as the response time of the video disc player exceeds one second when jumps have to be made of a thousand positions or more. However, this will seldom occur.

Although MUMPS is a language ideally suited for database development and retrieval, it cannot meet the speed of a compiler-based language as it is generally an interpreter. Every time MUMPS has to carry out a disc access, the block containing the desired information, is read from disc into central memory. A new block is only read from disc when the information required is not available in the blocks already in memory. The new block will replace the one that has not been referred to for the longest period. This block management is not controlled by the running computer program. However, when it is desirable to specify exactly which information must reside in central memory, this can be achieved by storing this information in local variables.

Retrieval speed can be improved in two ways:

- (1) The use of local variables to reduce the number of disc accesses.
- (2) The use of generated programs in AIDA, i.e. to change from table-driven to program-driven.

In the future we want to make use of graphics and multiple window scrolling to present information to the user. This requires a lot of calculations, which MUMPS is not very good at. Therefore a redesign of the user interface is being considered in the compiler-based language C. External function calls offer the possibility of integrating programs written in MUMPS and C, respectively, permitting us to exploit the characteristics of these two languages in the program.

Acknowledgments

The help of Prof. J.P.A. Baak in the project also providing clinical cases of ovarian pathology is gratefully acknowledged. The authors wish to thank B. Franken, P. van der Voort, R. van der Kooi, J. van der Lei and J.J. Luursema for their

assistance and advice. This work was supported by grant PF-1207 from the Praeventiefonds.

References

- [1] W.M. Boon, H.P. Westerhof, J.S. Duisterhout and P.V.M. Cromme, The role of AIDA in a primary care information system, *Comput. Methods Programs Biomed.* 25 (1987) 287-296.
- [2] N.J. Pressman, Archival systems in diagnostic quantitative pathology, in: *Proc. Int. Assoc. Pathol. Meeting, Vienna* (1986).
- [3] J.S. Duisterhout, B. Franken and F.S.C. Witte, Structure and software tools of AIDA, *Comput. Methods Programs Biomed.* 25 (1987) 259-274.
- [4] S.F. Serov, R.E. Scully and L.H. Sobin, Histological typing of ovarian tumours, in: *International Histological Classification of Tumours 9* (World Health Organisation, Geneva, 1973).
- [5] A. Blaustein, *Pathology of the Female Genital Tract*, 2nd edn. (Springer Verlag Berlin, 1982).
- [6] A.M. van Ginneken, Evaluation of Some Database Management Systems (Department of Medical Informatics, Free University, Amsterdam, 1986) (internal report).