

# GURU: Information Retrieval for Reuse

Yoëlle S. Maarek  
IBM, Haifa Research Laboratory  
Matam, Advanced Technology Center  
Haifa, 31905, Israel  
yoelle@haifa.vnet.ibm.com

Daniel M. Berry  
Technion, Israel Institute of Technology  
Computer Science Department  
Haifa, 32000, Israel  
dberry@csa.cs.technion.ac.il

Gail E. Kaiser  
Columbia University  
Department of Computer Science  
New York, NY 10027, USA  
kaiser@cs.columbia.edu

## Abstract

Although software reuse presents clear advantages for programmer productivity and code reliability, it is not practiced enough. One of the reasons for the only moderate success of reuse is the lack of software libraries that facilitate the actual locating and understanding of reusable components. This paper describes a technology for automatically assembling large software libraries that promote software reuse by helping the user locate the components closest to her/his needs.

Software libraries are automatically assembled from a set of unorganized components by using information retrieval techniques. The construction of the library is done in two steps. First, attributes are automatically extracted from natural language documentation by using a new free-text indexing scheme based on the notions of lexical affinities and quantity of information. Then, a hierarchy for browsing is automatically generated using a clustering technique that draws only on the information provided by the attributes. Thanks to the free-text indexing scheme, tools following this approach can accept free-style natural language queries.

This technology has been implemented in the GURU system, which has been applied to construct an organized library of AIX utilities. An experiment was conducted in order to evaluate the retrieval effectiveness of GURU as compared to INFOEXPLORER, a hypertext library system for AIX 3 on the IBM RS/6000 series, as well as to two other indexing schemes. We followed the usual evaluation procedure used in information retrieval, based upon recall and precision measures, and determined that our system retrieved more effectively than the others.

## 1 Introduction

Software reuse is widely believed to be a promising means for improving software productivity and reliability [Frakes and Gandel, 1989], and therefore is an issue of growing interest in software engineering. Unfortunately, not enough *adequate* libraries of reusable software components are available. By adequate, we mean that the library

- provides a sufficient number of components, over a spectrum of domains, that can be reused as is (*black-box* reuse) or easily adapted (*white-box* reuse), and
- is organized such that existing code closest to the users' needs is easy to locate. In particular, the library should provide mechanisms to help the reuser look for functionally close components that meet some given requirements.

This paper is concerned with the second adequacy issue, and more generally with library systems that provide means for representing, storing and retrieving reusable components.

The first stage in building a library consists of *indexing* the objects to be stored in it, that is, producing a set of characterizing attributes, or *profile*, for each of these objects. The profile for each object represents the reusable object. Therefore, the quality of indexing is crucial to the quality of the library. Functionality is an important aspect of software components. Thus, it is necessary to include conceptual information about functionality in the indices. Unfortunately, conceptual information is difficult to obtain. Few programmers provide conceptual indices for their code. Moreover, even if provided, they can hardly be expressed under a common formalism since pieces of code typically originate from multiple sources. One solution is to manually index software components *a posteriori* according to a given classifying scheme, but this task is expensive.

As an alternative, we propose to automatically identify indices by analyzing the natural-language documentation, in the form of manual pages or comments, usually associated with the code. Natural-language documentation is clearly a rich source of conceptual information. However, this information is contained only implicitly, in an unstructured way, and is not usable as such. In order to extract usable information from free-style documentation, we propose to use information retrieval techniques. Once the indices have been produced, components can be automatically classified, stored and retrieved according to their profiles.

The classifying stage in the construction of a library consists of gathering objects into classes such that the members of the same class share some set of properties. The basic motivation for classifying is to facilitate browsing among similar components in order to identify the best candidates for reuse, or at least, a set of potentially adaptable components. Browsing is more important for software libraries than for other kind of libraries, since there rarely exists a component perfectly matching a user's query. Moreover, local browsing allows the user to discover unanticipated opportunities for reuse.

We have designed and implemented a tool, GURU, that embodies the above approach. GURU automatically assembles conceptually structured software libraries from a set of unindexed and unorganized software components. In the first stage, GURU extracts the indices from the natural language documentation associated with the software components to be stored, by using a new indexing scheme. This indexing scheme is based on *lexical affinities* and on their statistical distribution. It identifies a set of attributes for each document to represent a functional description of the associated software unit. In the second stage, GURU assembles the indexed objects into a browsing hierarchy by using a *hierarchical clustering* technique that draws information exclusively from the indices identified in the previous stage. Thus, GURU supports both classical linear retrieval, in which candidates are ranked according to a numerical measure that evaluates how well they answer the query, and cluster-based retrieval, in which the browse hierarchy directs the search for the best candidate.

Section 2 briefly compares the artificial intelligence and information retrieval approaches to construction of software libraries and explains why we follow an IR approach. Section 3 describes the indexing method. Section 4 presents the classification approach and the clustering technique used for assembling the library. Section 5 deals with the retrieval stage. Section 6 gives results using our GURU implementation and a formal evaluation based on usual methodology for evaluating information retrieval systems. Finally, Section 7

summarizes the main contributions of this work. Related work is discussed as relevant throughout the paper.

## 2 AI vs IR approach

Previous efforts for building reuse systems can be roughly classified into two groups according to the approach<sup>1</sup> adopted, the free-text indexing approach as defined in information retrieval (IR), and the knowledge-based approach as defined in artificial intelligence (AI).

The IR free-text approach<sup>2</sup> consists of drawing information only from the structure of some documents that provide information about the software components. No semantic knowledge is used and no interpretation of the document is given. The reuse tool attempts to characterize the document rather than understand it. There are currently very few software library systems that follow such an approach, or use existing IR techniques. Among them, the RSL [Burton et al., 1987] system, for instance, automatically scans source code files and extracts comments explicitly labeled for reuse with attributes such as keyword, author, date created, etc. The keyword attribute provides a list of free-text single-term indices very much like those used in IR tools. The REUSE [Arnold and Stepoway, 1987] system provides a menu-driven front end to an information retrieval system. Thus all kinds of software objects, including user menus and system thesauri, are stored as textual documents. These two systems use some kind of IR related technique. However the only system, to our knowledge, that applies a pure IR free-text approach is the system proposed by Frakes and Nejme [Frakes and Nejme, 1987]. They use the CATALOG information retrieval system for storing and retrieving C software components. Each component is characterized by a set of single-term indices that are automatically extracted from the natural-language headers of C programs. Therefore, the construction of the C components repository is done automatically, and does not require any pre-encoded knowledge as in RSL for instance.

In contrast, in the knowledge-based approach, the reuse tool aims at understanding the queries and the functionality of components before providing an answer. Knowledge-based systems are often smarter than IR systems. Some of them are context sensitive and can generate answers adapted to the user's expertise. As a tradeoff, they require some domain analysis and a great deal of pre-encoded semantic information, which is usually provided manually. They are based upon a knowledge base that stores semantic information about the domain and about the language itself in case of a natural-language interface. The main problem of applying this approach in the context of software libraries is that many domains cannot be easily circumscribed and the domain analysis is very difficult [Devanbu, 1990]. This makes the construction of such systems very tedious and expensive. Examples of AI or knowledge-based reuse tools are numerous, *e.g.*, [Prieto Diaz and Freeman, 1987], [Wood and Sommerville, 1988], [Allen and Lee, 1989], [Devanbu et al., 1989], [Tichy et al., 1989].

The AI approach can be useful in some applications. However, we prefer the IR approach for reasons of:

- cost: the library system is built entirely automatically,

---

<sup>1</sup>Another approach is the hypertext approach (See [Frakes and Gandel, 1990] for a survey). We do not address this approach here, because we are concerned with the type of information used to *build* a library rather than with searching. The hypertext approach is orthogonal to the approaches described here; and hypertext tools can easily be integrated with most IR or AI-based reuse tools.

<sup>2</sup>For brevity, we will refer to this approach as the IR approach even though some IR techniques do not use free-text indexing.

- transportability: the library system can be rebuilt for any domain since it does not require manually provided domain knowledge,
- scalability: the repository can be easily updated when new components are inserted, either by re-compiling the indices or by applying incremental techniques, the indexing task is entirely mechanical.

We therefore apply a pure IR approach, in the same direction as Frakes and Nejme, by automatically building free-text indices that characterize software components. For more effective retrieval, we also use a free text method, which is richer than the single-term indexing used in the IR-based tools described in this section. The following section explains our source of information and how the indexing is performed.

### 3 The Indexing Stage

The major advantage of automatic indexing over manual indexing, besides the obvious cost considerations, is that it allows a unified scheme, insuring that indices will be compatible with each other. The idea is to extract attributes from an existing source of information, *i.e.*, the code and the natural-language documentation. Some work has been done towards extraction of primitive functional information from the code [Maarek and Kaiser, 1987], [Schwanke et al., 1989], however, the richer source of functional information is the natural-language documentation, assuming any is available.

An examination of numerous samples of code allowed us to reach the conclusion that some useful information can be extracted from programs written in a high-level language using good programming style, whereas little conceptual information can be found in typical real-world code chosen at random [Maarek, 1989]. Unfortunately, even when dealing with well-written code, there is a very low probability that the programming styles of the various pieces of code will be consistent. Even a single programmer may use totally different identifiers for expressing the same concept from one day to another. Since software components come from multiple sources in the context of large software libraries, extracting attributes from code would necessitate as many indexing schemes as there are code sources. Another limitation comes from the fact that there are many more possibilities for identifiers than for natural-language words since they do not follow any morphological or syntactic rules.

In other words, when there is no way to guarantee good, and let alone consistent and compatible, programming styles, extracting attributes from raw code does not give significant results. Therefore, we prefer concentrating on the other possible source of information, *i.e.*, the natural-language documentation either inserted into the code, *i.e.*, the comments, or associated with the code, *e.g.*, manual pages.

Comments are intended to help programmers understand the code and thus may provide functional information. They deal with specific parts of the code into which they are inserted, and they may give information on various parts at various levels of abstraction. Extracting functional information from comments entails two activities,

- defining an indexing scheme that allows extracting attributes from natural language phrases or sentences, and
- relating comments to the portion of code they concern.

The second activity is very complex in free-style code. Indeed, in free-style programming, programmers can insert comments wherever, and in any format and any length, they wish. Although comments usually describe the containing routine or the one just below, in general it is impossible to automatically determine

what part of the code is covered. A solution would be to consider that all the comments inserted in a specific piece of code constitute a global natural-language description of the considered code. Unfortunately, this is not the case. Comments rank from low-level implementation details to high-level description. For instance, in the `rm.c` source file in Berkeley UNIX, one can find comments as various as:

```
/* current pointer to end of path */ or
/* rm - for ReMoving files, directories & trees */
```

The first conveys no useful functional information while the second hits the mark exactly. In general, there are many more low level, and useless for our purpose, comments than high level ones, and there is no way to automatically distinguish between them. Therefore, so long as no style is enforced, it is very difficult to extract useful information from comments.

Let us note, however, that any piece of natural language, from comments inserted in the code to design specifications, which is specifically related to software code and whose level of abstraction is known can bring useful information. Thus, we have worked on extracting functional information from comments in the framework of RPDE [Harrison, 1986], a structured software development environment, in which comments are linked to the portion of code they describe. In the following, though, we try to remain as general as possible, and we do not assume that any commenting style is enforced. Therefore, although our indexing scheme is applicable to any piece of natural-language that brings some functional information, we will exemplify it through the analysis of manual pages clearly related to reusable components, such as UNIX-like manual pages.

In the rest of this paper, the AIX documentation is taken as our corpus since it fulfills the requirement of being structured into manual pages. Moreover the AIX documentation can be seen as a regular real-world documentation database since it is of average quality as far as commenting style is concerned. Many even consider the AIX documentation of poor quality when compared to Berkeley UNIX documentation due to typos, inconsistent style, poor vocabulary, etc.

### 3.1 A Richer Indexing Unit: the Lexical Affinity

There has been much work in IR dealing with natural-language text, a large variety of techniques have been devised for indexing, classifying and retrieving documents [Salton, 1989], [Salton and McGill, 1983]. One of the main concerns in IR is the automatic indexing of documents, which consists of producing for each document a set of indices that form a *profile* of the document. A profile is a short-form description of a document, easier to manipulate than the entire document, which plays the role of a surrogate at the retrieval stage.

Several issues need to be addressed when indexing a document with respect to the nature and the form of the produced indices. More precisely, the indexing vocabulary can be either controlled or uncontrolled. In the controlled vocabulary approach only a restricted set of indices are authorized (*e.g.*, in MEDLARS [Salton and McGill, 1983]), whereas in the uncontrolled vocabulary, or free text, approach, there is no constraint on the nature of the indices. It has been shown that both approaches are comparable in terms of performance, [Frakes and Gandel, 1989], [Salton and McGill, 1983], however we prefer the uncontrolled vocabulary approach in the context of software reuse, for the same reasons of cost, portability and scalability. Indeed, defining an adequate controlled vocabulary is a manual, domain-dependent task and, therefore, suffers from the same drawbacks as the encoding of a knowledge-base.

Another important issue in automatic indexing is the nature of the indices. The most usual form is a single-term index, each of which is a single word without contextual information. It has also been proposed to use term phrases as indexing units rather than single terms so as to refine the meaning of constituent

words. However, the use of word co-occurrences has not brought good results. As expressed by Salton [Salton, 1989] (p 296),

“... a phrase-formation process controlled only by word co-occurrences and the document frequencies of certain words is not likely to generate a large number of high-quality phrases.”

A possible solution to this problem is to use syntactic information, such as part-of-speech derived from specially formatted dictionaries [Klingbiel, 1973], in order to provide further control over phrase formation, or more refined analysis including semantics [Sparck Jones and Tait, 1984]. But,

“The available options in phrase generation appear limited, and the introduction of costly and refined methodologies may bring only marginal improvements.” [Salton, 1989] (p 298)

We are more optimistic, and believe that indexing units richer than single terms can be used and bring significant improvement at low cost. The atomic unit we propose to use in order to demonstrate this is derived from the notion of *lexical affinity*. In linguistics, a syntagmatic lexical affinity (LA), also termed *lexical relation*, between two units of language stands for a correlation of their common appearance in the utterances of the language [de Saussure, 1949]. The observation of LAs in large textual corpora has been shown to convey information on both syntactic and semantic levels, and provides us with a powerful way of taking context into account [Smadja, 1989].

We propose to use the notion of LA for indexing purposes, and restrict the above definition by observing LAs within a finite document rather than within the whole language so as to retrieve *conceptual* affinities that characterize the document, rather than purely *lexical* ones. Moreover, we only consider LAs involving *open-class words* as meaning-bearing, whereas LAs involving *closed-class* words<sup>3</sup> are not.

Ideally, LAs are extracted from a text by parsing it since two words share a lexical affinity if they are involved in a modifier-modified relation. Unfortunately, automatic syntactic parsing of free-style text is still not very efficient [Salton and Smith, 1989]. Instead, we make use of simple co-occurrence. It has been shown by Martin *et al.* that 98% of lexical relations relate words that are separated by at most five words within a single sentence [Martin et al., 1983]. Therefore, most of the LAs involving a word  $w$  can be extracted by examining the neighborhood of each occurrence of  $w$  within a span of five words (-5 words and +5 words around  $w$ ).

The extraction technique consists of sliding a window over the text and storing pairs of words involving the head of the window (if it is an open-class word) and any of the other open-class elements of the window. The window is slid word by word from the first word of the sentence to the last, the size of the window decreasing at the end of the sentence so as not to cross sentence boundaries<sup>4</sup>, since lexical affinities cannot relate words belonging to different sentences. The window size being smaller than a constant, the extraction of LAs is linear in the number of words in the document. An algorithm for the sliding window technique is presented in Figure 1. Maarek and Smadja have used a similar technique in [Maarek and Smadja, 1989], which was also based on Martin’s results [Martin et al., 1983], but more adapted to the analysis of large corpora.

In summary, the first stage in indexing a manual page consists of extracting all the potential LAs by using the sliding window technique. Once extracted, the potential LAs are stored under their canonical form, in which each word is represented by its inflectional root, or lemma, *i.e.*, the singular form for nouns

---

<sup>3</sup>In general, open-class words include nouns, verbs, adjectives and adverbs, while closed-class words are pronouns, prepositions, conjunctions and interjections.

<sup>4</sup>The isolation of sentences, based on sentence-ending punctuation, is the only parsing performed.

```

For each sentence  $S$  in the document  $d$ 
  For each word  $w$  in  $S$  from the beginning to the end of  $S$ 
     $w \leftarrow \text{lemma}(w)$ 
    (where  $\text{lemma}(w)$  represents the inflectional root of  $w$ )
  EndFor
  For each lemma  $w$  in  $S$  from the beginning to the end of  $S$ 
    If  $w$  is an open-class word then
      Let  $w_1, \dots, w_n$  be the  $n$  words immediately following  $w$  in  $S$ 
      (where  $n = 5$  except when the end of the sentence is reached earlier)
      For  $i = 1$  to  $n$ 
        If  $w_i$  is an open-class word then
          Get  $f$ , frequency count of  $\{w, w_i\}$ 
          ( $f = 0$  when the LA has not been encountered before)
          Store  $\{w, w_i\}$  with a frequency count of  $f + 1$ 
        EndIf
      EndFor
    EndIf
  EndFor
EndFor

```

Figure 1: Sliding window technique

and the infinitive form for verbs. The potential LAs extracted from the manual page of `mv` in AIX and ranked by frequency of occurrence are presented in Table 1. For the sake of the comparison, a list of the single words extracted from the same manual page is shown in the first column, also ranked by frequency of appearance.

open-class words	freq	LAs	freq
file	30	file move	9
directory	14	be file	8
mv	11	directory file	7
files	8	file system	5
new	7	file overwrite	5
name	7	file mv	5
move	7	file name	4
newname	6	name path	3
is	6	do file	3
system	5	directory move	3
one	5	different file	3
...	...	...	...

Table 1: Keywords and lexical affinities classified by frequency in the `mv` manual page

Among the extracted lexical relations, some correspond to abstractions of the considered document, and some do not. Since we are interested in indexing textual documents, in the first stage, we isolate actual affinities by using frequency criteria. It has been demonstrated that the frequency of occurrence of a term within a document is related to the importance of the word in a text [Luhn, 1958]. This is also true for the common appearance of pairs of words and even more for lexical affinities.

### 3.2 From LAs to Indices

When analyzing a document, many potential lexical affinities are thus identified. Some of these lexical affinities are conceptually important and some are not. As seen in Table 1, frequency of appearance is a good indicator of relevance. However, some noise exists, mainly due to words appearing too often in a given context. In order to reduce the influence of such words, it is necessary in the second stage to select from among the lexical affinities identified only the most representative ones, *i.e.*, those containing the most information.

We have defined a measure evaluating the *resolving power* of an LA. It is based upon the quantity of information of each of the words involved in the LA, as well as upon the frequency of appearance of this LA within the considered document. The *quantity of information* of a word within a corpus is defined as:

$$\text{INFO}(w) = -\log_2(P\{w\}) \quad (1)$$

where  $P\{w\}$  is the observed probability of occurrence  $w$  in the corpus [Ash, 1965], [Salton and McGill, 1983]. Therefore, the more frequent a word is in a domain, the less information it carries. From this definition, we infer the definition of the quantity of information of an LA  $(w_1, w_2)$  as:

$$\text{INFO}((w_1, w_2)) = -\log_2(P\{w_1, w_2\}) \quad (2)$$

To simplify the computation of this factor, in the rest of this work, we consider words within the textual universe as independent variables<sup>5</sup>. Thus, we use the following formula for computing the quantity of information of an LA.

$$\text{INFO}((w_1, w_2)) = -\log_2(P\{w_1\} \times P\{w_2\}) \quad (3)$$

Then, we define the resolving power of an LA in a given document as follows. Let  $(w_1, w_2, f)$  be a tuple retrieved while analyzing a document  $d$ , where  $(w_1, w_2)$  is an LA appearing  $f$  times in  $d$ . The *resolving power*<sup>6</sup> of this LA in  $d$  is defined as:

$$\rho((w_1, w_2, f)) = f \times \text{INFO}((w_1, w_2)) \quad (4)$$

The higher the resolving power of a lexical affinity, the more characteristic of the document it is. The resolving power allows us to evaluate the importance of a lexical affinity within a text by taking into account both its frequency of appearance in the text and the quantity of information of the words involved. Thus, even though the lexical affinity **(file system)** appears as often as **(file overwrite)** in the **mv** manual page, it has only a lower resolving power, simply because the word “system” has a lower quantity of information than “overwrite” in the AIX documentation.

In order to be able to compare the relative performances of different documents, in terms of resolving power, we transform the raw  $\rho$  score into a standardized score. The standardized score, or z-score, is defined as  $\rho_z = (\rho - \bar{\rho})/\sigma$  where  $\bar{\rho}$  and  $\sigma$  are the average and standard deviation of the  $\rho$ -values. This transformation does not alter the distribution and allows us to evaluate the relative status of the score in the  $\rho$  distribution. In the rest of this paper, the  $\rho$ -values we give as examples will therefore represent the z-score rather than the raw score.

Table 2 compares the list of LAs for the **mv** manual page ranked by frequency and by resolving power. In it, the LA *(file move)* has a greater resolving power than any of the following LAs. Moreover, some noisy LAs such as *(do file)* or *(be file)* (in italic fonts in the table) have disappeared because *both* words involved in the LAs are highly frequent in the corpus and thus have a low quantity of information.

<sup>5</sup>This assumption represents only an approximation since words in English are definitely not independent, but are distributed according to the rules of the language.

<sup>6</sup>This notion is related to that of mutual information [Ash, 1965].

LAs	freq	LAs	$\rho$
file move	9	file move	8.38
<i>be file</i>	8	file mv	4.36
directory file	7	directory file	4.03
file system	5	file overwrite	3.87
file overwrite	5	directory move	1.98
file mv	5	file system	1.95
file name	4	mv rename	1.71
name path	3	move mv	1.58
<i>do file</i>	3	different file	1.40
directory move	3	name path	1.33

Table 2: Comparison of frequency and  $\rho$ -value for the LAs in **mv**

For each document, we select as indices those LAs with the highest resolving power. More precisely, we are interested in the LAs that represent peaks in the distribution of  $\rho$ -values. Therefore, we keep as indices only the LAs whose  $\rho$  value is one standard deviation above the mean, *i.e.*, such that  $\rho \geq \bar{\rho} + \sigma$ , where  $\bar{\rho}$  represents the mean and  $\sigma$  the standard deviation of the distribution of  $\rho$  values within one document. The choice of such a threshold<sup>7</sup> is reflected in Tables 2, 3 and 4, where only LAs with a z-score greater than 1 are presented.

LAs	$\rho_z$
copy file	6.49
directory file	2.47
file source	2.15
infile subdirectory	1.98
contain subdirectory	1.30
copy cp	1.30
copy regular	1.02

Table 3: LAs ranked by  $\rho$ -values for **cp**

LAs	$\rho$
directory make	5.08
create mkdir	2.74
directory mkdir	2.55
directory permission	1.48
directory write	1.03

Table 4: LAs ranked by  $\rho$ -values for **mkdir**

The set of LAs of a document selected by ranking  $\rho$ -values and taking those one standard deviation above the mean forms the profile of the document. The major contribution of this technique consisted in adapting the notion of lexical affinity for indexing purposes. We gave some intuitive indications on how an LA-based indexing scheme is richer than a single-word scheme. We will demonstrate later that it ensures a better retrieval effectiveness.

<sup>7</sup>This classical threshold guarantees to keep only a small percentage of the sample elements in most distributions.

The next section explains how software components can be stored and classified using the profiles produced at the indexing stage.

## 4 The Classifying Stage

Normally, when a user wants to use a software library, s/he first has to access a library that might contain the desired component, then has to provide a formal description of the researched component according to the vocabulary understood by the library system. Unfortunately, in most cases, this ideal scenario does not work out. The main reason is that in real life applications, the component perfectly matching the user's requirements does not exist in the library, or it is not indexed as the user had guessed it would be.

In such cases, a traditional database management system fails to help the user. Indeed, to be retrieved from the database, a component must exactly match the query<sup>8</sup>. Such strict matching is inappropriate in a software library system since the user often cannot know the exact characteristics of the desirable component and, even when s/he does, there is rarely a perfect match.

Software libraries should not only permit retrieving candidate components that perfectly or partially match the query, but also permit browsing among components that share some functionality. It is therefore desirable to structure the library for making the search, retrieval and browsing mechanisms as fast and convenient as possible, in order to make the access to the library attractive.

We propose here to perform the search and retrieval operations using a conventional inverted index file structure, and to cluster the library in order to facilitate the browsing operation. Section 4.1 explains how the index repository is built using an inverted file structure, and Section 4.2 presents the clustering technique used to build the browse hierarchy. Section 5 explains how they are used to perform the search and browsing operations.

### 4.1 Building the index repository

The information about the indices is stored in the form of a standard inverted file index, in which every index identified in the collection of documents points towards the list of documents in which it appears. More precisely, each index points to a list of *postings*, where each posting contains not only a document identifier but also some additional information (see Figure 2).

In our application, indices take the form of LA's as well as single words. Index LA's are defined as tuples  $(w, w')$  in which  $w$  precedes  $w'$  in the lexicographic order. The reason for ordering  $w$  and  $w'$  is to avoid duplicate LA's by forcing every LA into a canonical form. Moreover, we also store  $w$  and  $w'$  as individual indices in order to detect partial matching, *i.e.*, only one word in common between query LA's and document LA's, as fast as possible.

Every index points towards a list of postings of the form  $(d, \rho)$ , in which  $d$  is the document whose profile contains the index and  $\rho$  is its corresponding normalized resolving power. The normalization performed is the *cosine normalization* [Salton and McGill, 1983], which consists of dividing each index score by the norm of the document in a vectorial space. The norm is defined as the square root of the sum of the scores

---

<sup>8</sup>A notable exception is ARES [Ichikawa and Hirakawa, 1986], a relational database that allows flexible interpretation of queries. In ARES the similarity between elements can be evaluated via a lookup in a table that has to be provided beforehand. ARES is not discussed here since its purpose is not to classify software. Further, it has the drawback of requiring a great deal of pre-encoded knowledge.

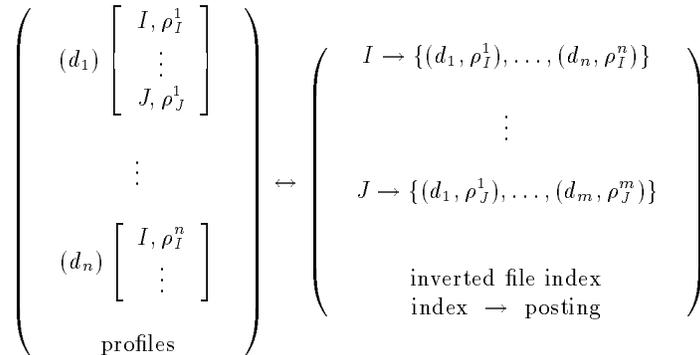


Figure 2: From profiles to inverted file index

of all the indices of the profile. More formally, let  $\rho_i$  be the resolving power of the  $i^{th}$  index in a given document  $d$ , the norm of the document is defined as:

$$|d| = \sqrt{\sum_{i=1, n} \rho_i^2}$$

In order to have fast access to the postings, the index part of the inverted file index (also called the dictionary) is usually implemented via B-trees, Patricia tries, or hash tables.

The stored information is used to retrieve and rank candidates as explained in Section 5.

## 4.2 Building the browse hierarchy

Browsing is crucial in software library systems. The most common way to make browsing operations possible is to group items judged to be similar by using clustering operations. It has been pointed out that “associations between documents convey information about the relevance of documents to requests” [Jardine and van Rijsbergen, 1971]. Thus, many clustering methods have been used for information retrieval [Can and Ozkarahan, 1983], [Griffiths et al., 1984]. The most popular clustering methods are the hierarchical agglomerative clustering (HAC) methods because their search and construction techniques are more efficient than for most non-hierarchical methods.

The following sections define some terminology in cluster analysis, describe the algorithms we used to build the browse hierarchy, and present some samples of the browsing hierarchy obtained for the AIX library.

### 4.2.1 Some terminology in cluster analysis

Classification by cluster analysis has been of long-standing interest in statistics as well as various other fields. It can be traced back to the work of Adanson in 1757 [Adanson, 1757], who used numerical clustering for classifying botanic species. Statisticians and taxonomists have widely developed the field since then.

Cluster analysis now offers a wide range of techniques for identifying underlying structures in large sets of objects and revealing links between objects or classes of objects. One particular application of classification is the building of libraries.

There is no strict definition of cluster, but it is generally agreed that a cluster is a group of objects whose members are more similar to each other than to the members of any other group. Typically, the goal of cluster analysis is to determine a set of clusters, or a clustering, such that inter-cluster similarity is low and intra-cluster similarity is high. The similarity between objects is evaluated via a numerical measure called a *dissimilarity index* defined as follows.

**Definition 1** *Let  $\Omega$  be a set of objects. A dissimilarity index  $\delta$  over  $\Omega^2$  is a function from  $\Omega \times \Omega$  to  $R_+$  that satisfies the following properties,*

$$(i) \quad \forall o \in \Omega, \delta(o, o) = 0, \tag{5}$$

$$(ii) \quad \forall (o, o') \in \Omega^2, \delta(o, o') = \delta(o', o). \tag{6}$$

*Note that a distance is a dissimilarity index but that a dissimilarity index does not necessarily satisfy the triangle inequality and therefore is not a distance.*

The dissimilarity index between objects is used as the basic criterion to determine clusters. Clustering techniques allow identifying not only clusters but also relationships among them. The structure of the set of clusters as well as their internal structure vary with the clustering technique. Clustering methods are usually classified<sup>9</sup> according to the structure of the set of clusters produced, *e.g.* hierarchical, flat, overlapping, etc., as well as the technique used, *e.g.*, divisive, agglomerative, incremental, etc. As explained previously, hierarchical agglomerative techniques are very convenient for building browse hierarchies. The basic principle that these techniques follow is presented below.

Hierarchical numerical clustering aims at building hierarchies, over a set of objects, in which each internal node corresponds to a cluster of objects and each leaf represents an individual object, or more precisely a singleton cluster. Most hierarchical clustering methods are based upon the same general method, called the Hierarchical Agglomerative Clustering (HAC) method, which consists of iteratively gathering objects into clusters, until only one cluster remains.

The HAC general method iteratively builds a sequence of partitions or *level clusterings* of  $\Omega$ , that is, a sequence of disjoint clusters covering the original set of objects,  $\Omega$ . The level clusterings form coarser and coarser partitions by an iterative process, beginning with the level clustering formed by the set of singletons in the power set  $\wp(\Omega)$ , *i.e.*,  $\{\{o_1\}, \{o_2\}, \dots, \{o_n\}\}$ , and ending up with the coarsest partition of  $\Omega$ , *i.e.*,  $\{\Omega\}$ . The final output of this clustering process is a particular form of hierarchy called a *dendogram*. The HAC general method can be expressed as follows:

- **Start** with the subset of  $\wp(\Omega)$  formed by singleton elements.
- **Repeat** the following steps iteratively **until** there is only one cluster.
  - **Identify** the two clusters that are the most similar.
  - **Merge** them together into a single cluster.

---

<sup>9</sup>With the recent introduction of conceptual clustering [Michalski and Stepp, 1983], another distinction has been introduced according to the definition of the clusters obtained, in extension (*i.e.*, by enumeration of its members) for regular (or numerical) clustering and in intension (*i.e.*, by membership rules) as well as in extension for conceptual clustering.

The HAC method requires a measure of similarity not only over the set of objects, but also over the set of clusters. The dissimilarity index between clusters is usually derived from a user-given dissimilarity index,  $\delta$ , between objects. The way of defining  $\Delta$  has a direct influence on the final form of the hierarchy obtained. Once a dissimilarity index  $\delta$  between objects is provided, HAC methods differ only by the choice of this measure. The most commonly used HAC methods are the single link and complete link methods [Lance and Williams, 1967]. Many other methods such as the centroid method, Ward’s method, etc., define still other dissimilarity indices but most of them require the dissimilarity index over  $\Omega$  to be a distance, that is, to satisfy the triangle inequality. The reader should consult [Everitt, 1980] for an extensive survey of the HAC methods. The time complexity of the HAC algorithm is at most  $O(n^2 \log n)$  where  $n$  is the number of objects involved. For some particular definitions of  $\Delta$ , it can be reduced to  $O(n^2)$ .

#### 4.2.2 Adapting a clustering technique for building a browse hierarchy

As explained above, we propose to use a HAC technique to generate a browse hierarchy. In this perspective, we (1) need to define a measure of similarity between the objects considered, *e.g.*, the documents, and (2) explain how to make a browse hierarchy out of the dendrogram generated by the HAC technique. Let us address these two points.

In information retrieval, numerous measures of similarity between documents, also termed measures of association or coefficients of association, have been defined [van Rijsbergen, 1979]. The simplest of all is defined as:

$$|X \cap Y| \tag{7}$$

where  $X$  and  $Y$  are the profiles of two documents. This measure represents the number of common index units. Various other measures [van Rijsbergen, 1979] have been defined such as:

$$\frac{2|X \cap Y|}{|X| + |Y|} \quad \text{Dice's coefficient} \tag{8}$$

$$\frac{|X \cap Y|}{|X \cup Y|} \quad \text{Jaccard's coefficient} \tag{9}$$

$$\frac{|X \cap Y|}{|X| \times |Y|} \quad \text{Salton's Cosine coefficient} \tag{10}$$

They can all be considered as normalized versions of (7) since they are functions of the cardinality of  $X$ ,  $Y$ ,  $X \cap Y$ , or  $X \cup Y$ .

In our context, we have more information than just the presence or absence of index units in the profile, and therefore we propose to take into account the  $\rho$ -values of LAs in the evaluation of the measure of association between documents. For any profile  $X = \{(i, \rho)\}$ ,  $p(X)$  is the projection set of  $X$  over the set of indices. Then, the simplest measure is  $|p(X) \cap p(Y)|$ , *i.e.*, the number of indices in common in both profiles. In order to take into account the resolving power of indices as well, we define our measure  $\delta$  for two profiles  $X$  and  $Y$ , such that  $X \neq Y$ , as

$$\delta(X, Y) = \sum_{i \in p(X) \cap p(Y)} (\rho_X(i) \times \rho_Y(i)) \tag{11}$$

where  $\rho_X(i)$  is the  $\rho$  value of the index  $i$  in the profile  $X$ , and similarly for  $Y$ . Note that  $\delta$  is a measure of similarity rather than a measure of dissimilarity. Its inverse is a measure of dissimilarity as long as  $\delta(X, X)$  is set to a sufficiently large arbitrary value so that its inverse can be considered essentially null.

Given such a measure of similarity between profiles, we define a measure of similarity between clusters, according to the single link or complete link techniques for instance, and then use the hierarchical

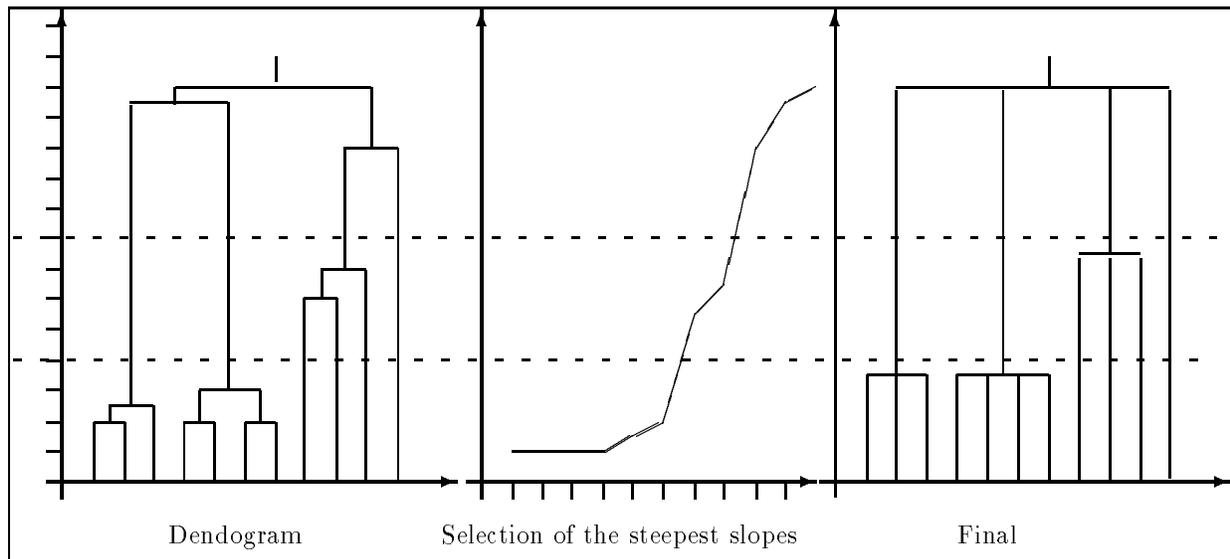


Figure 3: Principle of selection of level clusterings

agglomerative clustering algorithm in order to build a browse hierarchy of software components. We also did some experiments in earlier versions of GURU using an incremental conceptual clustering technique [Maarek, 1990] for constructing the browse hierarchy. However, despite interesting results, the cost of building and maintaining the hierarchy was prohibitive (exponential time like for most conceptual clustering techniques) when compared to regular clustering techniques and did not appear to be better in terms of retrieval effectiveness.

Every HAC techniques builds a binary hierarchy. Not all levels of the hierarchy are equally significant; therefore, the usual approach is to select manually the most significant level clusterings, this task being usually performed by a data analyst. The following proposes a method for automatically identifying the most useful level clusterings, thus producing a binary hierarchy that is not necessarily binary.

This method of selection is based on the following principle. Each level clustering in the dendrogram corresponds to the merging of two clusters in the previous level clustering and therefore to a particular value of the similarity measure. If we label the dendrogram with these values  $y_n, \dots, y_1$ ,  $n$  being the number of objects, from the bottom to the top of the hierarchy, it can easily be shown that the  $y_i$ 's are non-strictly monotonically increasing, for dissimilarity measures, or decreasing, for similarity measures, for the single and complete link clustering methods. We select those levels that correspond to the gap in the distribution of  $y_i$ 's by (1) plotting the segment connecting the pairs  $y_{i+1}, y_i$  from  $i = n - 1$  to  $i = 1$ , and (2) keeping the levels that correspond to the steepest slopes. This represents the intuitive method that a data analyst would apply. Figure 3 gives an intuitive presentation of the method *via* an example whereas Figure 4 gives the formal algorithm. The time complexity of the latter is linear in the number of objects.

### 4.3 Some examples

Portions of the browse hierarchy built from the AIX documentation are shown in Figures 5 and 6. In Figure 5, some interesting clusters are isolated. Thus, in the figure we have a cluster gathering commands related to the manipulation of regular expressions, and a cluster gathering editors. These two clusters are

```

Let  $y_n, \dots, y_1$  be the merging values of the similarity measure from the bottom to the top.
For  $i = n - 1$  to  $i = 1$ 
     $\Delta y_i = y_{i+1} - y_i$ 
    (evaluate the slope of the connecting segment)
EndFor
Compute  $\overline{\Delta y}$  the mean of the  $y_i$ 's
Compute  $\sigma$  the standard deviation of the  $y_i$ 's around  $\overline{\Delta y}$ 
Let  $t(k) = \overline{\Delta y} + k\sigma$ 
    (where  $t(k)$  corresponds to a threshold defined by  $k \geq 0$ )
For  $i = n - 1$  to  $i = 1$ 
    If  $\Delta y_i > t(k)$ 
        Select level clustering  $i$ 
    EndIf
EndFor

```

Figure 4: Selection of level clusterings

also part of the same super-cluster, mainly because these editors require manipulating regular expressions. Then, there are two outliers that could not be included in a cluster: `makekey` and `termdef`. Then a small cluster groups `ps` and `kill`, which both are strongly related as they give information about processes or handle them. Finally, there are two big clusters, one for yellow pages commands and another for SCCS routines. The clustering is not always of such good quality as can be seen in Figure 6, either because of the nature of the documentation or because of the principle of clustering itself. For instance, the commands `xcalc` and `dc`, which both are calculators, belong to the same cluster, but `bc` has been forgotten in this cluster. This is due to the fact that the manual page of `bc` does not refer to the concept of calculator at all, but defines `bc` as an interpreter for an arithmetic language. The real problem with clustering is illustrated with the third cluster in this figure, which gathers `batch`, `at`, `crontab`, `date` and `istat`. This cluster has been formed because all these commands are related to the notion of date or time; unfortunately, this is not the main functionality of all of these commands and therefore this cluster is somehow misleading. Note, however, that the lower level cluster including `at` and `batch` is a good one. Also while the cluster does not gather components that have similar functionality, it does gather components that are constructed from the same library of date and time functions. One of the authors actually thought that the clustering to be quite natural and useful. He had developed a Hebrew date program by invoking the `date` program and a public domain Gregorian-Hebrew date converter from inside a quick and dirty program that converts the output of the first to the form required by the second. This is a classic example of reuse to save work for the programmer.

The hierarchy thus generated is used as an aid to browse when nothing relevant has been retrieved via linear retrieval, or in order to increase recall since there is no way to be sure that all the relevant components have been retrieved at the linear retrieval stage. It can also be used as the basic repository to be searched during retrieval, but we prefer to use the traditional linear retrieval technique instead because it is clearly more trustworthy considering the problems described above.

By nature, this indexing technique suffers from noise since it is based on only statistical observations. Noisy indices involve generally misspelled or non meaningful strings of characters that are mixed with natural language (for describing instructions for instance), or side concepts such as the time, day and month in the example cited above. This noise cannot be avoided when dealing with free-style text.

Fortunately, these noisy LAs do not cause real trouble at the linear retrieval stage since there is a very

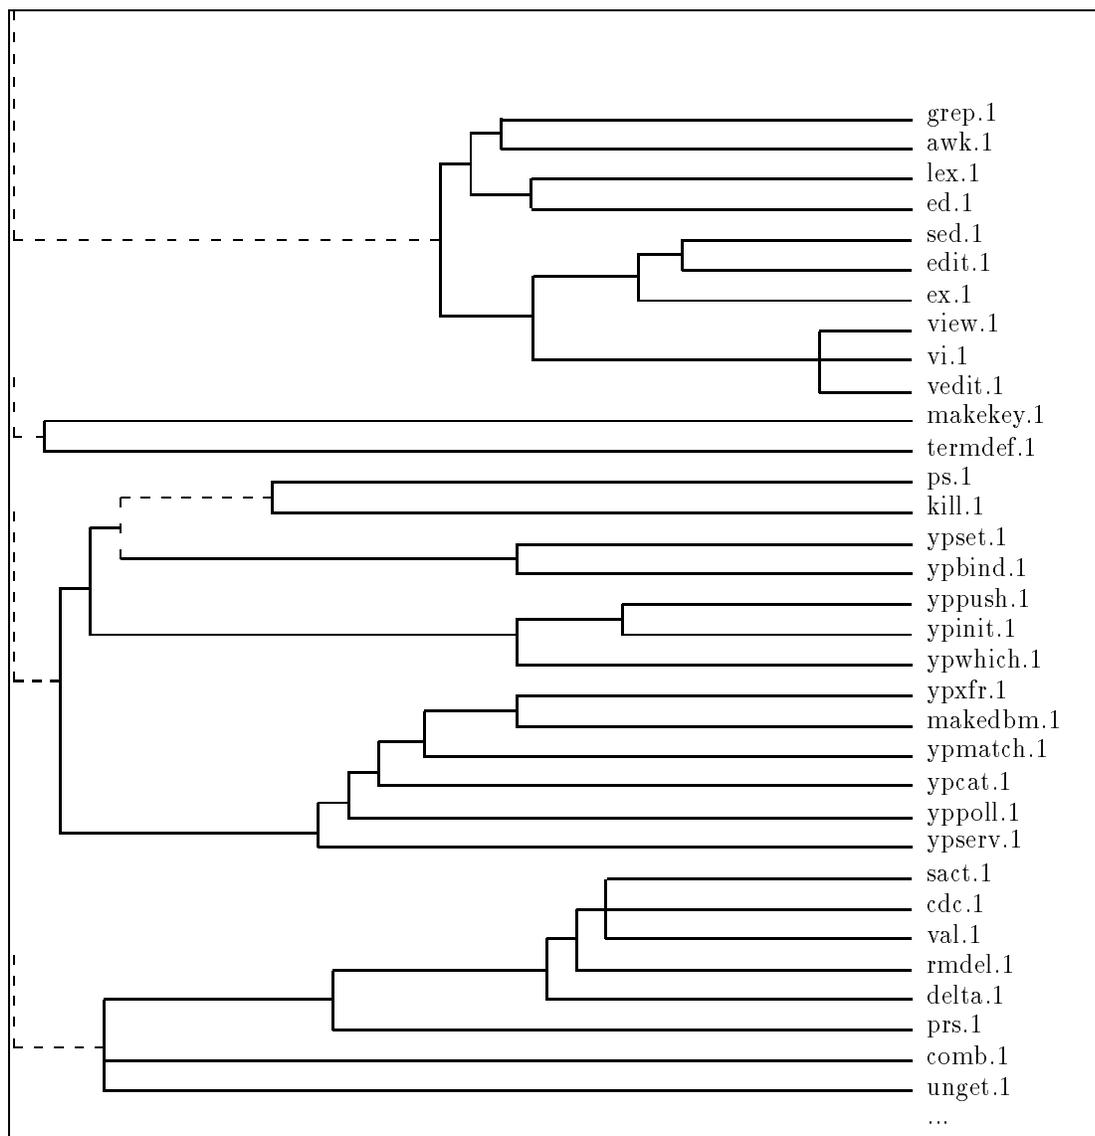


Figure 5: Portion of AIX hierarchy (single link,  $k=0.5$ )

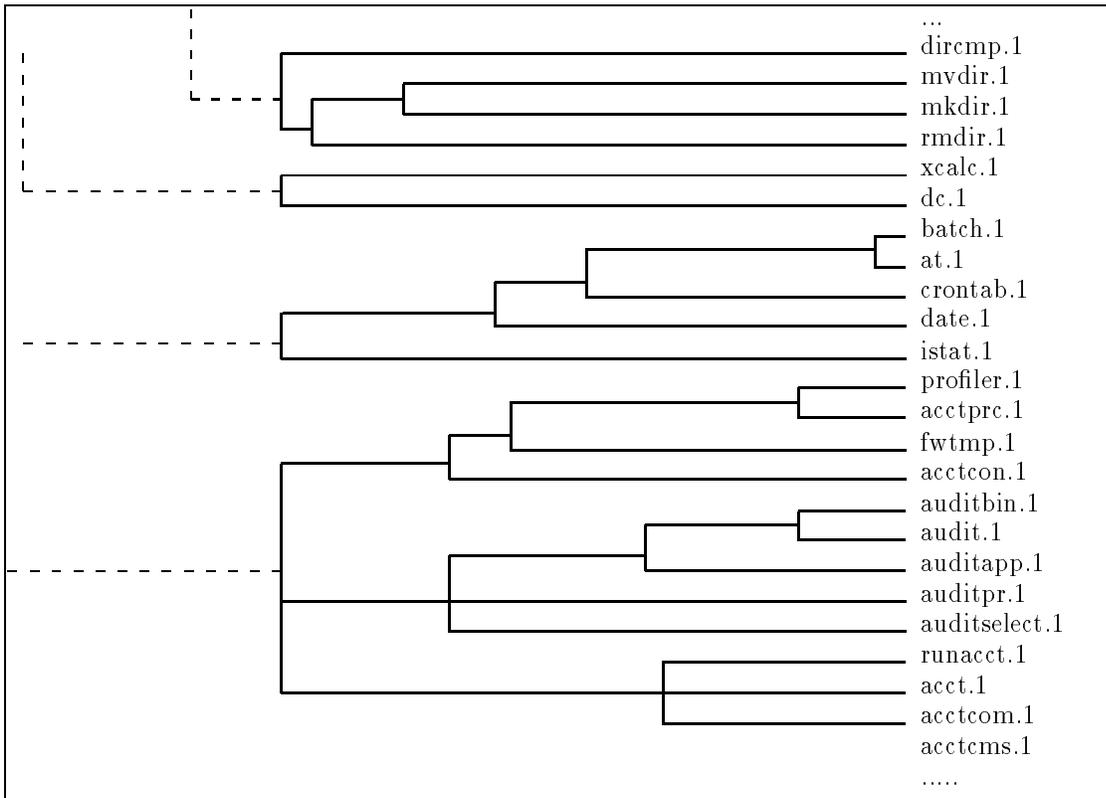


Figure 6: Portion of AIX hierarchy (single link,  $k=0.5$ )

low probability that the user would use non meaningful character strings in her/his queries. So noisy LAs are part of the profiles of components but rarely lead to the selection of the considered component. On the other hand, noisy LAs might induce the formation of poor quality clusters, but generally only higher levels of the hierarchy are affected since side concepts are not given much weight when evaluating similarity. Section 5.3 explains how this browsing hierarchy is used at retrieval stage.

## 5 The Retrieval Stage

The previous sections explain how libraries of reusable components are assembled. We also need to be able to retrieve the components that match the requirements when at least one exists, or to assist in the selection of the closest components via a browsing facility.

The usual scenario when retrieving a component is the following:

- **Query specification:** The user expresses a query according to the authorized vocabulary.
- **Linear retrieval:** A search locates the candidate components and the candidates are ranked according to their degree of match with the query.
- **Browsing** Cluster-based retrieval is initiated when adequate components have not been found by the linear retrieval.

The following explains how these three stages are supported in our approach.

### 5.1 Query specification

Using uncontrolled-vocabulary indexing, as we do, presents clear advantages at the query specification stage. Indeed, a minimum of constraint is put on the user as s/he expresses her/his query. The user does not have to learn a specific index language or understand the organization of the library. In this approach, queries are usually either boolean or free-style. In the boolean model, queries consist of a set of terms connected by boolean operators such as AND, OR, or NOT. Long boolean queries, especially when adjacency operators are allowed, are difficult to formulate. Even a computer literate might have trouble for instance to express the query: “human factors and/or system performance in medical database” [Harman, 1992] as a boolean expression. Moreover, users might have to reformulate their query several times (changing some ORs into ANDs, or vice versa), if too many or not enough candidates are returned the first time. For user friendliness, free-style queries are clearly preferable queries. Section 6.1.2 describes the problems we encountered while experimenting with a tool that implements boolean queries.

In GURU, the user can express her/his query in natural language and then the indexing component is applied in order to translate the query into attributes understandable by the system. Exactly the same technique is used for extracting LAs from natural-language queries as from natural-language documentation. This provides a very convenient and user-friendly interface between the user and the library system, because the user is not constrained by any rigid formalism.

The queries can be expressed in free-style natural language. However, the user must be aware of the fact that queries are not really interpreted, but rather considered as a description of the functionality of the desired component. For instance, the user could express queries of the form “How can I do such and such”

since only the “such and such” would be considered for indexing, the rest being either closed-class words or words with low quantity of information. Formulating a query that necessitates some understanding, such as a query including negations like “but not”, would lead only to wrong interpretations.

## 5.2 Linear retrieval

In order to retrieve the best candidates for a given query, we apply the usual IR method, which consists of considering the query as a document and retrieving the components in the repository whose profiles are the most similar to the profile of the query. A possible measure of similarity is the  $\delta$  measure defined in (11), Section 4.2.2. The retrieval engine needs to store in an accumulator the degree of similarity of each component with the query. A time efficient method of computing these scores is presented in Figure 7. Those components with the highest scores are then returned to the user, ranked in order of decreasing similarity with the query.

```

Initialize the score accumulator  $a$ , which stores the degree of similarity of each document
with the query.
Index the query  $q$ , using the same indexing mechanism as used for documents
For each term  $t_j$  in the query's profile  $\{(t_1, w_1^q), \dots, (t_j, w_j^q)\}$ 
    Fetch the postings list  $p_j$  of  $t_j$  in the inverted file index of the considered collection
    For each posting  $(d_i, w_j^i)$  in  $p_j$ 
        Compute the new score of document  $d_i$  in the accumulator as  $a_i \leftarrow a_i + (w_j^q \times w_j^i)$ 
    EndFor
EndFor
Sort the documents according to the value of their score in the accumulator  $a$ 
Return the list of document ranked by decreasing score

```

Figure 7: Linear retrieval technique

In Figure 8, a session example with GURU is presented. This example shows the advantage of analyzing the entire manual page rather than only the most significant sections such as the “Purpose” or “Description” sections. As the reader can see, there is only a minimal overlap between the query and the “Purpose” section of the `xsend` manual page. Note that the first retrieved candidate (with the highest score) is perfectly relevant to the query. In the software library application, unlike other applications such as legal or medical applications, it is expected that only a few candidates are highly relevant to the query. Thus the retrieval algorithm needs to give the relevant candidates a distinctively high score. This is the case in our example, in which the right candidate has the highest score. The user does not need then to scan the subsequent candidates which have much lower scores as they are only slightly related to the query.

## 5.3 Browsing, cluster-based retrieval

The retrieval stage, in classical library management systems, is often limited to locating a set of components exactly matching the user’s query or, when such components do not exist, related components. Library systems do not usually provide any further assistance.

In our approach, the user may communicate interactively with the system in order to direct the browsing when the first retrieval yield is not satisfactory. The linear search retrieves the most related candidates,

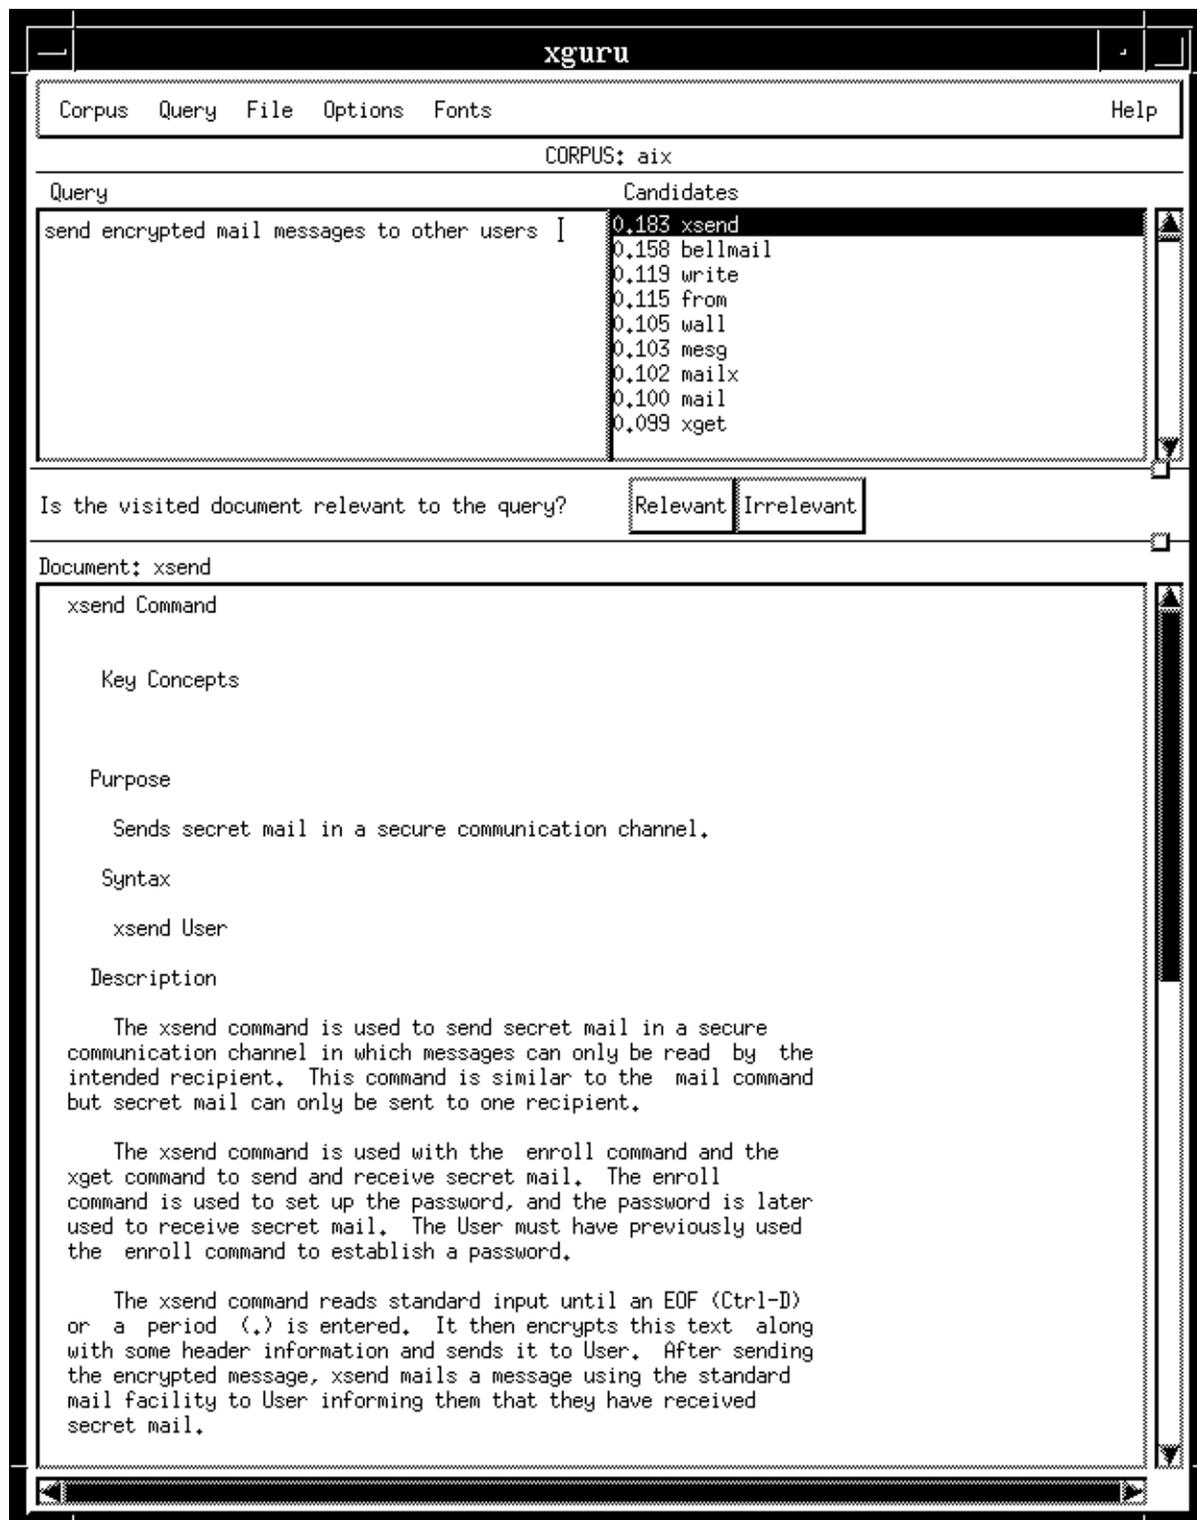


Figure 8: Session example with GURU

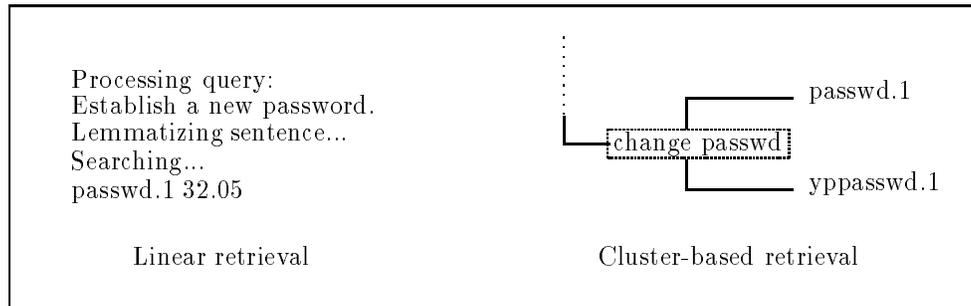


Figure 9: Browsing in the hierarchy

and then the browsing process begins.

Typically, the user starts from one of the candidates retrieved by the linear search and explores the hierarchy bottom-up. Consider the browse hierarchy given in Figure 5 and suppose that a user gives a query asking about ways “to identify a process”. If the first candidate retrieved at the retrieval search is **kill**. Then, the user can access the browse hierarchy, and explore the clusters including **kill** in order to determine which components are strongly related. In our example, s/he will find component; **ps** is clearly a better candidate for this given query than the one retrieved by the linear search. Another example is illustrated in Figure 9. The two relevant candidates in AIX for the query “establish a new password” are **passwd** and **yppasswd**. However, the linear retrieval retrieves only **passwd** simply because the query had no intersection with the profile of **yppasswd**. At this point, the user could reformulate the query, but s/he might not be aware that s/he has missed some relevant candidates. Using the browse hierarchy is therefore more convenient in order to check if some unexpected candidates have been missed. In the example, both **passwd** and **yppasswd** are strongly related: their profiles share the LA (*change passwd*)<sup>10</sup>, and therefore belong to the same low-level cluster in the browsing hierarchy. Browsing in the hierarchy from **passwd** allows the user to retrieve the other relevant candidate. These two examples show how a browse hierarchy can help improve the finding of possible candidates that could be missed via linear retrieval.

At any point, the user can consult the profile of a component in order to have more information about its functionality. Fast access to profiles is achieved via the profile repository. The user can also provide, at any stage, further information in order to get a finer retrieval. By browsing, s/he gets more information about components and learns how to provide discriminating queries.

## 6 Empirical Results

The approach described in the previous sections has been embodied in a tool, GURU, which has been fully implemented, in C, under AIX on an RS/6000. GURU comes with a Motif and an emacs interfaces. The system is currently used at several IBM sites by a growing pool of users, and yields quality results.

We have tested our system on the entire AIX documentation available to us, which describes approximately 1,100 AIX components. When building the index repository, we therefore processed the entire documentation that forms a corpus of more than 800,000 words, and we identified 18,000 LAs for the

<sup>10</sup>Note that “passwd” here is proper name and is different from the noun “password” mentioned in the query.

1,100 profiles.

In order to evaluate GURU's performance, we used the following criteria.

- **User effort.** This consists of all the effort that must be expended by the user in order to use the library system. It is impossible to formally measure user effort. However, thanks to the uncontrolled vocabulary approach that we applied, we believe that the effort that must be invested for using GURU is minimal. Queries can be formulated in natural language, and therefore the user is not required to learn any index language and formalism.
- **Maintenance effort.** This consists of all the effort that is necessary to keep the system working and up to date. This effort includes, in particular, indexing new components and adding them to the library. The maintenance stage is highly facilitated in GURU. The indexing is performed automatically and the insertion of new components can be done incrementally. Kaplan and Maarek, in [Kaplan and Maarek, 1990], have described several algorithms for incrementally updating a repository of LA-based indices when inserting, deleting or modifying components. The first author is currently integrating into GURU other incremental schemes which are not described here.
- **Efficiency.** This refers to the average interval between the time a query is issued and the time an answer is given. Efficiency becomes an issue only if a retrieval takes so long that users start to complain. Our experience with the system shows that efficiency is not an issue, as the response time is reasonable. Profiling the execution of time to perform the query was dominated by the time to map the repository file into the address space of the query program. The lookup operations and the printing of the LA-file name pairs consumed almost no time in comparison. Test queries involving from 5 to 15 LAs each took approximately 2.5 seconds on an RT, and 0.15 seconds on an IBM RISC System/6000. The better performance of the latter is partly due to its more efficient implementation of file mapping.
- **Retrieval effectiveness.** This is clearly the most important performance criterion. It refers to the system's ability to provide information services as needed by the user.

The next section focuses on evaluating the retrieval effectiveness of GURU.

## 6.1 Measuring Retrieval Effectiveness

### 6.1.1 Recall and Precision

*Recall* and *precision* [Salton and McGill, 1983] are the most widely used measures for evaluating retrieval effectiveness. Recall and precision can be defined intuitively as “the whole truth and nothing but the truth”. More formally, recall is defined as the proportion of *relevant* material, *i.e.*, it measures how well the considered system retrieves *all* the relevant components. Precision is defined as the proportion of retrieved material that is relevant, *i.e.*, it measures how well the system retrieves *only* the relevant components. Recall can also be interpreted as the probability that a relevant component will be retrieved, and precision as the probability that a retrieved component will be relevant [Blair and Maron, 1985].

Recall and precision can be defined also as follows. Let  $C$  be the whole collection of components forming the library. For each query,  $C$  can be partitioned into two disjoint sets,  $R$ , the set of relevant material and  $\bar{R}$  the set of irrelevant material. Given the query, the system retrieves a set of components  $c$  that can also

be partitioned into relevant and irrelevant material, respectively,  $r$  and  $\bar{r}$ . Recall and precision are defined as:

$$\text{recall} = \frac{r}{R} \tag{12}$$

$$\text{precision} = \frac{r}{c} \tag{13}$$

Measurement of recall and precision require the ability to distinguish between relevant and irrelevant material. Relevance judgements are always debatable since they depend on the test collection on which the measurements are performed; nevertheless they provide a good indication of retrieval effectiveness. There exist numerous standard IR test collections; however, because of our domain of application, we chose a test collection especially designed for software libraries: the “AIX test collection” [Maarek, 1991].

### 6.1.2 Experiments and Comparison

This section describes the experiments that allowed us to evaluate the retrieval effectiveness of GURU. As a basis for comparison, we have considered INFOEXPLORER. INFOEXPLORER is a commercial CD-Rom Hypertext Information Base Library [InfoExplorer, 1990] for the IBM RS/6000 series that gives access to the documentation for AIX and to associated programs. Besides regular hypertext facilities, INFOEXPLORER provides a search facility that represents a good reference for comparison purposes since it is specifically designed for AIX. This search facility implements boolean queries with frequency ranking, *i.e.*, a model different from GURU’s. This approach, which is inspired by the database approach, is one of the most common in commercial help systems, but yet very primitive in terms of state-of-the-art IR. Therefore, we also compared, on the same test collection, GURU’s indexing scheme to other IR indexing schemes to show the contribution of LA’s and of our resolving power measure to the effectiveness of GURU.

Those schemes are presented below:

- INFOEXPLORER. Key features: single words, frequency-based indexing scheme, boolean queries. Its search facility allows a compound search, which is actually a search via boolean queries. The query is expressed as a word pattern formed of single words related by three possible connectors, “and”, “or”, and “butnot”. Moreover, users can restrict their search. They can give constraints specifying whether the pattern words must appear within the same article or within the same paragraph, the proximity of these words within a paragraph, and the search fields and the search categories. When given such a query, INFOEXPLORER returns a list of candidates that exactly fit the query, ranked according to the frequency of the pattern in the considered document.
- SCHEME 1. Key features: single word indices, entropy based score, free-text queries. We modified GURU, so as to consider only single words as indexing units. The purpose of using this scheme as a basis of comparison was to judge the influence of LA’s on the retrieval effectiveness.
- SCHEME 2. Key features: single word indices,  $tf \times idf$  score, free-text queries. This scheme implements the classical “tfidf” score [Salton and McGill, 1983], which is defined as follows. Let  $i$  be a single word index in a document  $d$ , its associated score is  $tf \times idf$  where  $tf$  represents the term frequency of  $i$  in  $d$ , and  $idf$  the inverse document frequency of  $i$ , *i.e.*, the number of documents in the collection that contains the term  $i$  at least once.

The test collection free-text queries could not be fed directly to INFOEXPLORER since the latter’s compound search facility accepts only boolean queries. Therefore, feeding the queries to INFOEXPLORER

required some supplementary effort, first choosing the right connectors between open-class words extracted from the queries, and possibly dropping some words when the recall was too low. In our interaction with the compound search facility, we had to refine and retry the query formulation several times. We kept only the best result for comparison purposes, since we wanted to compare the tools' indexing schemes rather than their querying facilities.

The comparison was performed by measuring, for both systems, precision at several levels of recall. We followed the usual procedure [van Rijsbergen, 1979], [Salton and McGill, 1983], which consists of:

1. plotting precision-recall curves for each test query with each plot corresponding to a given cutoff value,
2. extrapolating these curves so as to obtain precision values for recall values that were not effectively achieved, and finally,
3. deriving from the curves computed at the previous stage the average precision values at fixed recall intervals, so as to obtain a single average precision recall curve for the system considered.

We have built such curves for GURU, INFOEXPLORER, and the two other schemes and plotted them on the same axes (see Figure 10). The best performance is reached by the system whose curve is closest to the area where both precision and recall are maximized, the upper right corner of the graph. As mentioned, because of the indexing scheme of both systems, we could expect that INFOEXPLORER would achieve a lower precision but higher recall than GURU. It turned out that the maximum recall achieved by both systems was approximately the same, all ranks included, but, from the graph presented in Figure 10, it is clear that GURU has better precision than INFOEXPLORER. In other words, GURU achieves a higher precision without losing in recall. Note that INFOEXPLORER's effectiveness, in spite of its primitive indexing scheme, had better results than the single word-tfidf scheme. This is due to the fact that we only kept the best results among those obtained from various boolean translations of a same query. What was gained from the effectiveness perspective was actually lost from the user-friendliness perspective.

These results show that for the sample tested, GURU achieves higher precision than INFOEXPLORER without losing recall. For this sample, the recall rate is increased when we make use of the GURU browse facility. For instance, in several cases, some related components were not retrieved during linear retrieval, but only during browsing.

## 7 Conclusion

We have presented a method for automatically constructing software libraries from a collection of documented but unindexed software components. We discussed the advantages of using natural-language documentation as opposed to source code as a source of functional information. We then described a new free-text indexing scheme, for automatically producing document profiles, based upon a richer unit than single terms, the lexical affinity. All associated software components could then be classified, stored, compared and retrieved, via linear or cluster-based techniques, according to these indices.

These methods and schemes are embodied in a new tool which has been implemented and evaluated for retrieval effectiveness. The evaluation compared GURU with the INFOEXPLORER hypertext library, built specifically to help find software components in the AIX system, as well as with more advanced IR indexing schemes. The average recall-precision curves of all tools were computed. The results of this test indicate that GURU's effectiveness was better than the other ones.

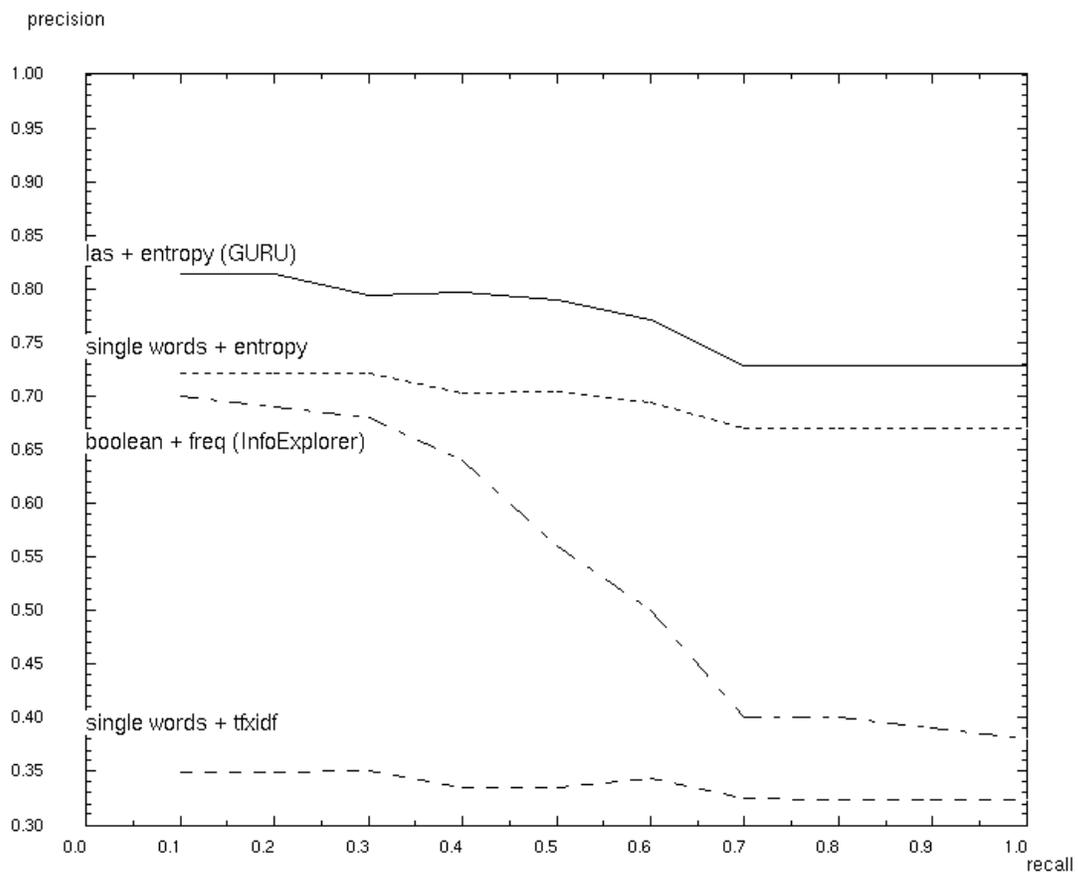


Figure 10: Precision-recall curves (means across queries)

The major contribution of this work consists of bringing classical and new information retrieval techniques to bear in software reuse. This involved

- designing a new indexing scheme based on high information content lexical affinities,
- adapting classical numerical cluster analysis techniques for assembling software components into browse hierarchies, and
- designing retrieval mechanisms specifically adapted to the LA-based indexing scheme so as to provide a complete storage and retrieval framework.

Finally, the evaluation we have performed seems to indicate that Salton's statement about the limitation of the phrase generation approach in indexing (See Section 3.1) is overly pessimistic and that significant improvements over single terms techniques can be achieved at relatively low cost.

## Acknowledgments

An earlier version of this paper appears as "An Information Retrieval Approach for Automatically Constructing Software Libraries", *IEEE Transactions on Software Engineering*. 17:8, pp 800-813, August 1991.

When the preliminary part of this work was done, Prof. G. Kaiser was supported by National Science Foundation grants CDA-8920080, CCR-8858029 and CCR-8802741, by grants from AT&T, BNR, Citicorp, DEC, IBM, Siemens, Sun and Xerox, by the Center for Advanced Technology and by the Center for Telecommunications Research.

## References

- [Adanson, 1757] Adanson, M. (1757). *Histoire Naturelle du Sénégal. Coquillages. Avec la relation abrégée d'un voyage fait en ce pays, pendant les années 1749, 50, 51, 52 et 53*. Bauche, Paris, France.
- [Allen and Lee, 1989] Allen, B. and Lee, S. (1989). A knowledge-based environment for the development of software parts composition systems. In *Proceedings of the 11<sup>th</sup> ICSE*, pages 104–112, Pittsburgh, PA.
- [Arnold and Stepoway, 1987] Arnold, S. and Stepoway, S. (1987). The reuse system: Cataloging and retrieval of reusable software. In Tracz, W., editor, *Software Reuse: Emerging Technology*, pages 138–141. Computer Society Press.
- [Ash, 1965] Ash, R. (1965). *Information Theory*. Interscience Publishers (John Wiley & Sons), New York.
- [Blair and Maron, 1985] Blair, D. and Maron, M. (1985). An evaluation of retrieval effectiveness for a full-text document retrieval system. *Communications of the ACM*, 28(3):289–299.
- [Burton et al., 1987] Burton, B., Aragon, R. W., Bailey, S., Koelher, K., and Mayes, L. (1987). The reusable software library. In Tracz, W., editor, *Software Reuse: Emerging Technology*, pages 129–137. Computer Society Press.
- [Can and Ozkarahan, 1983] Can, F. and Ozkarahan, E. (1983). A clustering scheme. In *Proceedings of SIGIR'83*, pages 115–121, Bethesda, MD. ACM Press.
- [de Saussure, 1949] de Saussure, F. (1949). *Cours de Linguistique Générale*. Librairie Payot, Paris, France, quatrième edition.
- [Devanbu, 1990] Devanbu, P. (1990). Re-use of software knowledge: A progress report. In *Third Annual Workshop: Methods and Tools for Reuse*, Syracuse, NY.

- [Devanbu et al., 1989] Devanbu, P., Selfridge, P., Ballard, B., and Brachman, R. (1989). A knowledge-based software information system. In *Proceedings of IJCAI'89*, pages 110–115, Detroit, MI.
- [Everitt, 1980] Everitt, B. (1980). *Cluster Analysis*. Halsted Press (John Wiley & Sons), New York.
- [Frakes and Gandel, 1989] Frakes, W. and Gandel, P. (1989). Classification, storage and retrieval of reusable components. In Belkin, N. and van Rijsbergen, C., editors, *Proceedings of SIGIR'89*, pages 251–254, Cambridge, MA. ACM Press.
- [Frakes and Gandel, 1990] Frakes, W. and Gandel, P. (1990). Representing reusable software. *Information and Software Technology*.
- [Frakes and Nejme, 1987] Frakes, W. and Nejme, B. (1987). Software reuse through information retrieval. In *Proceedings of the 20th Annual HICSS*, pages 530–535, Kona, HI.
- [Griffiths et al., 1984] Griffiths, A., Robinson, L., and Willett, P. (1984). Hierarchical agglomerative clustering methods for automatic document classification. *Journal of Documentation*, 40(3):175–205.
- [Harman, 1992] Harman, D. (1992). Ranking algorithms. In Frakes, W. B. and Baeza-Yates, R., editors, *Information Retrieval, Data Structure and Algorithms*, pages 241–263. Prentice Hall.
- [Harrison, 1986] Harrison, W. (1986). A program development environment for programming by refinement and reuse. In *Proceedings of the 19th HICSS*, pages 459–469, Kona, HI. CS Press.
- [Ichikawa and Hirakawa, 1986] Ichikawa, T. and Hirakawa, M. (1986). Ares: A relational database with the capability of performing flexible interpretation of queries. *IEEE Transactions of Software Engineering*, 12(5):624–634.
- [InfoExplorer, 1990] InfoExplorer (1990). *IBM AIX Version 3 for RISC System/6000. Commands Reference*. IBM, 1st edition.
- [Jardine and van Rijsbergen, 1971] Jardine, N. and van Rijsbergen, C. (1971). The use of hierarchic clustering in information retrieval. *Information Storage and Retrieval*, 7(5):217–240.
- [Kaplan and Maarek, 1990] Kaplan, S. and Maarek, Y. (1990). Incremental maintenance of semantic links in dynamically changing hypertext systems. *Interacting with Computers*, 2(3).
- [Klingbiel, 1973] Klingbiel, P. (1973). Machine-aided indexing of technical literature. *Information Storage and Retrieval*, 9:79–84.
- [Lance and Williams, 1967] Lance, G. and Williams, W. (1967). A general theory of classificatory sorting strategies. *Computer Journal*, 9:373–380.
- [Luhn, 1958] Luhn, M. (1958). The automatic creation of literature abstracts. *IBM Journal of Research and Development*, 2(2):159–165.
- [Maarek, 1989] Maarek, Y. (1989). *Using Structural Information for Managing Very Large Software Systems*. PhD thesis, Technion, Israel Institute of Technology, Haifa, Israel.
- [Maarek, 1990] Maarek, Y. (1990). An incremental conceptual clustering algorithm with input-ordering bias correction. In Golombic, M., editor, *Advances in Artificial Intelligence, Natural Language and Knowledge Base Systems*. Springer Verlag.
- [Maarek, 1991] Maarek, Y. (1991). Software library construction from an IR perspective. *SIGIR Forum*, 25(2):8–18.
- [Maarek and Kaiser, 1987] Maarek, Y. and Kaiser, G. (1987). On the use of conceptual clustering for classifying reusable ada code. In *Ada Letters, Using Ada: ACM SIGAda International Conference*, pages 208–215, Boston, MA. ACM Press.
- [Maarek and Smadja, 1989] Maarek, Y. and Smadja, F. (1989). Full text indexing based on lexical relations. an application: Software libraries. In Belkin, N. and van Rijsbergen, C., editors, *Proceedings of SIGIR'89*, pages 198–206, Cambridge, MA. ACM Press.
- [Martin et al., 1983] Martin, W., Al, B., and van Sterkenburg, P. (1983). On the processing of a text corpus: From textual data to lexicographic information. In Hartmann, R., editor, *Lexicography: Principles and Practice*, London. Applied Language Studies Series, Academic Press.
- [Michalski and Stepp, 1983] Michalski, R. and Stepp, R. (1983). Automated constructions of classifications: Conceptual clustering versus numerical taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 5(4):396–409.

- [Prieto Diaz and Freeman, 1987] Prieto Diaz, R. and Freeman, P. (1987). Classifying software for reusability. *IEEE Software*, 4(1):6–16.
- [Salton, 1989] Salton, G. (1989). *Automatic text processing, the transformation, analysis and retrieval of information by computer*. Addison-Wesley, Reading, MA.
- [Salton and McGill, 1983] Salton, G. and McGill, M. (1983). *Introduction to Modern Information Retrieval*. Computer Series. McGraw-Hill, New York.
- [Salton and Smith, 1989] Salton, G. and Smith, M. (1989). On the application of syntactic methodologies in automatic text analysis. In *Proceedings of SIGIR'89*, pages 137–150, Cambridge, MA. ACM Press.
- [Schwanke et al., 1989] Schwanke, R., Altucher, R., and Platoff, M. (1989). Discovering, visualizing and controlling software structure. In *Proceedings of the Fifth International Workshop on Software Specifications and Design*, pages 147–150, Pittsburgh, PA.
- [Smadja, 1989] Smadja, F. (1989). Lexical co-occurrence: The missing link. *Journal of the Association for Literary and Linguistic Computing*. Oxford University Press, 4(3).
- [Sparck Jones and Tait, 1984] Sparck Jones, K. and Tait, J. (1984). Automatic search variant generation. *Journal of Documentation*, 40(1):50–66.
- [Tichy et al., 1989] Tichy, W., Adams, R., and Holter, L. (1989). NLH/E: A natural-language help system. In *Proceedings of the 11<sup>th</sup> ICSE*, pages 364–374, Pittsburgh, PA.
- [van Rijsbergen, 1979] van Rijsbergen, C. (1979). *Information Retrieval*. Butterworths, second edition.
- [Wood and Sommerville, 1988] Wood, M. and Sommerville, I. (1988). An information retrieval system for software components. *SIGIR Forum*, 22(3,4):11–25.