

**Legal Knowledge Based Systems**  
**JURIX'96**  
**Foundations of legal knowledge systems**

The Foundation for Legal Knowledge Systems

Editors:  
R.W. van Kralingen  
H.J. van den Herik  
J.E.J. Prins  
M. Sergot  
J. Zeleznikow

Word use in legal texts: statistical facts and practical applicability 91-100  
*Kees van Noortwijk and Richard V. De Mulder*

The publisher is aware that, in spite of all care taken by editors, authors and publisher, shortcomings may occur in this publication. For these, however, editors, authors and publisher cannot accept any liability.

Paper version: Tilburg University Press, ISBN 90-361-9657-4, NUGI 699

1996 JURIX The Foundation for Legal Knowledge Systems

Alle rechten voorbehouden. Niets uit deze uitgave mag worden verveelvoudigd, opgeslagen in een geautomatiseerd gegevensbestand, of openbaar gemaakt, in enige vorm of op enige wijze, hetzij elektronisch, mechanisch, door fotokopieën, opnamen, of op enige andere manier, zonder voorafgaande schriftelijke toestemming van de stichting JURIX.

Voor zover het maken van kopieën uit deze uitgave is toegestaan op grond van artikel 16b Auteurswet 1912 j° het Besluit van 20 juni 1974, Stb. 351, zoals gewijzigd bij het Besluit van 23 augustus 1985, Stb. 471 en artikel 17 Auteurswet 1912, dient men de daarvoor wettelijk verschuldigde vergoedingen te voldoen aan de Stichting Reprorecht (Postbus 882, 1180 AW Amstelveen). Voor het overnemen van gedeelte(n) uit deze uitgave in bloemlezingen, readers en andere compilatiewerken (artikel 16 Auteurswet 1912) dient men zich tot de stichting JURIX te wenden.

## CONTENTS

Deep models, ontologies and legal knowledge based systems <i>T.J.M. Bench-Capon and P.R.S. Visser</i>	3
The formal specification of a legal ontology <i>Pepijn Visser and Trevor Bench-Capon</i>	15
An integrated view on rules and principles <i>Bart Verheij</i>	25
A generic model for the interpretation of vague norms <i>Jeannette Quast, Jaap van den Herik and Leo Aarts</i>	39
Salomon: automatic abstracting of legal cases for effective access to court decisions <i>Caroline Uyttendaele, Marie-Francine Moens, Jos Dumortier</i>	47
A task-based hyperindex for legal databases <i>Luuk Matthijssen</i>	59
A framework for self-explaining legal documents <i>L. Karl Branting and James C. Lester</i>	77
Word use in legal texts: statistical facts and practical applicability <i>Kees van Noortwijk and Richard V. De Mulder</i>	91
Automating legal reasoning in discretionary domains <i>Andrew Stranieri and John Zeleznikow</i>	101
Legal reasoning by structural analogy based on goal dependent abstraction <i>Tokuyasu Kakuta, Makoto Haraguchi and Yoshiaki Okubo</i>	111
Using set-of-support control strategy to deal with indeterminacy in legal reasoning <i>Bipin Indurkha</i>	123
Constructing Normative Rules <i>Nienke den Haan and Joost Breuker</i>	135

**WORD USE IN LEGAL TEXTS:  
STATISTICAL FACTS AND PRACTICAL APPLICABILITY**

Kees van Noortwijk and Richard V. De Mulder  
*Centre for Computers and Law*  
*Erasmus University*  
*P.O. box 1738, 3000 DR Rotterdam, The Netherlands*  
*email: {vannoortwijk, demulder}@rrj.frg.eur.nl*

**Summary**

This article starts by describing a research project in which the word use in two types of legal documents, namely statute law texts and case law texts, is compared to the word use in general Dutch texts. Quantitative methods were used to make these comparisons. The data resulting from this research can be used in several ways, two of which are described here in the form of applications which have been realised. The first application is capable of forming clusters of documents in a legal database. The second application makes it possible to define legal concepts which can be used to search a legal database.

**1 Introduction**

It goes without saying that language plays an essential role in the field of law. Statute law and Case law are expressed in language, as are the opinions of legal writers and colleague lawyers. An important part of a lawyer's work consists of interpreting such sources. After he or she has formed an opinion it is again by means of language that this opinion is communicated to others. This could of course be done orally, but most legal statements are probably made - or also made - in written form.

In this article, we will concentrate on legal language in written form. We will not take the perspective of the lawyer who *uses* or *produces* certain legal texts as part of his daily work. Instead we will look at these texts from the perspective of the legal researcher. We want to learn as much as possible about them. As we work with actual material, and try to gather reliable information on the language lawyers use, what we describe here can be characterised as *empirical legal research*, often referred to by the term *Jurimetrics research*<sup>1</sup>.

A common characteristic of most publications on legal language is that they concentrate on *semantic* aspects, *i.e.*, aspects which have to do with the *meaning* of the text. An example is research which concentrates on ascertaining which objectives of a certain piece of statute law are the most important. This is often connected to questions of a pragmatic kind, for example if these objectives could be fulfilled in practice. In addition, attention is sometimes paid to the *form* legal texts have, but in most cases this is studied in relation to the meaning. For example: because of long or unusual words or because of a difficult sentence structure legal texts are difficult to understand.

However, studying the form of legal texts can be interesting without dealing with the meaning of the text as well. Research of this kind is described in Van Noortwijk (1995). Some examples taken from this thesis are given in the next paragraph. It is found, for instance, that the word use in statute law texts differs measurably from that in case law texts, and that both differ even more from general Dutch texts. Having obtained detailed knowledge on this subject, an interesting question is, of course, if this knowledge can be put into use for the legal field. This question is addressed at the end of the second and in the third section of this article.

---

<sup>1</sup> See for instance Loevinger (1949), Kerkmeester (1989) and De Mulder (1984).

## 2 Studying the word use in legal texts

### 2.1 Some important terms

When a human reader, even one who does not have any legal knowledge, compares statute law or case law texts to other texts, it is quite likely that he will notice differences in form. For example, more formal words are used, sentences are structured in a different way, and certain headers are used. In many cases these differences are so obvious that it is possible to distinguish the legal from a general text type just by glancing at them for a few seconds. This being the case, an important question from the scientific point of view is if these differences in form can in some way or another be made concrete and expressed in terms of measurable figures. To address this question we have used some methods from the field of *quantitative linguistics* (sometimes also referred to as *statistical linguistics*)<sup>2</sup>. This is a branch of linguistic science in which the measurability of linguistic phenomena plays a central role. We will limit ourselves to the *word use* in texts for the moment and ignore characteristics such as the structure and the length of sentences or longer text entities.

Doing research into the word use in certain text types used to be very labour-intensive. To get a good overview, it is often necessary to study large text files or *corpora*. Word lists of these corpora had to be compiled manually until the early 1960's. From then, computers gradually took over this task. Nowadays, large amounts of text material is available in electronic form because of "electronic publishing" and electronic typesetting. Fast computers are available to almost every researcher. Word counting and syntactic analysis of text material can now be done in a relatively short time. This makes it possible to compile and compare very large text corpora, for instance a corpus containing every piece of statute law which is in force in a certain country at a certain moment. Some limitations still exist. Even at the syntactic or form level it is sometimes desirable to be able to understand the *meaning* of a text. However, in most cases computers are not capable of this, and therefore are unable to render the meaning of the words they count. Consequently, some characteristics of a word are difficult to record, such as the lexical category to which it belongs. We have avoided this limitation by using nothing but the original word forms from the text corpora. This means that we have not tried to reduce a word to its "stem" form (lemmatisation) or to distinguish different lexical categories.

### 2.2 Methods for analysing word use

Quantitative linguistic provides a number of methods to analyse the word use in a corpus. Characteristics which play an important role in these methods are for instance *word frequency* (how often does a certain word appear), *frequency distribution* (what is the pattern of the word frequencies of all the different words in a corpus) and *distribution of word types* (is a certain word used in every document in a corpus, or only in a subset of documents). These characteristics can be analysed by compiling a *frequency list* of the corpus. This is a list of all the different words (or "word types") in the corpus, plus the number of times the word appears in the corpus and the number of documents of which it is a part.<sup>3</sup> This list is sorted according to word frequency, the most common word being at the top. Based on this frequency list a number of linguistic measurements can be made, such as the "characteristic *K*" of Yule/Herdan<sup>4</sup>. The value of these measurements provides a typology of what could be called the "structure of word use". Apart from these data, which characterise the *way in which words are used*, the words themselves have also been studied in this research project. Points which have been taken into account in

<sup>2</sup> See for instance Guiraud (1959), Herdan (1966). For an overview of the developments in the field of quantitative linguistics, see Bailey (1969) and (more recently) Baayen (1989).

<sup>3</sup> See for examples of the use of these characteristics for instance Kucera and Francis (1967).

<sup>4</sup> Van Noortwijk (1995), p. 27.

this respect are, for instance, *word lengths* and the specific words which appear at the “head” of the frequency list (the most common words in a corpus).

### 2.3 Legal language versus “general Dutch”

The characteristics mentioned in the former section can be used to record certain quantitative data or measurements about a text corpus. To apply these characteristics, however, a suitable corpus must first be available. In the research project described here we have chosen to work with an extensive selection of legal and general documents. Two legal corpora were compiled. The first one contained every piece of Dutch statute law which was in force at the moment the corpus was created. The second one contained a broad selection of case law, namely all cases which were published in the leading magazine “Nederlandse Jurisprudentie” (Dutch Case Law) between 1965 and 1989. Together these two corpora contain well over 50 million words (or *word tokens*, in linguistic terms). To compare the word use in legal documents with that in other text types a third corpus was created, containing “general Dutch” text material. The material for this corpus was obtained from the Institute of Dutch Lexicology. It included texts from novels, short stories, text books and articles on all kinds of subjects.

For these three corpora, the quantitative characteristics mentioned in the former section were recorded by means of a series of computer programs designed especially for this purpose. As a first step in this process, three separate frequency lists were compiled, as well as lists of word types sorted according to their length and other necessary data. A comparison of the characteristics of the three corpora yielded the following similarities and dissimilarities.

- The general Dutch corpus on the whole contains a relatively high number of different words. Conversely, a word type is repeated more often in the legal texts than in the general texts.
- By plotting frequency data graphically and using statistical techniques such as regression analysis it is possible to distinguish a certain “head” in the frequency lists. This first part of the lists, containing the most common words, is *shorter* in both legal corpora than in the general Dutch corpus.
- A number of words from the top part of the frequency list clearly have *higher frequencies* than words at the same position on the list of the general corpus. Together with the former point, this means that in legal texts there is a *smaller* “core” vocabulary, of which every word type is used *more often*.
- The same conclusion can be drawn from a comparison of the *frequency distributions*. Another fact which emerges from these distributions is that although the general Dutch corpus contains a far greater number of different words, many of these words (a higher percentage than in the legal corpora) have very low frequencies (they appear less than ten times).
- The distribution of words over the documents of the corpus is somewhat less in the legal corpora than in the general corpus. This seems to indicate that many of the words used in statute law or case law texts are less “universal”, their use being limited to certain documents.
- Word use in each of the corpora can be effectively characterised by several “linguistic constants”. The value of the characteristic *K* which we mentioned earlier for instance ranges from 0.0128 for the statute law corpus, via 0.0111 for the case law corpus, to 0.0106 for the general Dutch corpus. These values are practically the same for each of the corpora as a whole, and for subsets from them.
- When *word lengths* are compared, it is noticeable that especially statute law texts contain a larger variety of different words with a length of 4, 5 or 6 characters. This appears to be caused by a greater number of actual *numbers* in these text (numbers of articles, sections, etc.).

- When we look at the words which are used most often, again very clear differences become visible. Of course certain special terms, like “artikel” (article), “wet” (law), “gemeente” (local council) etc. are used more often in statute law texts, while terms like “beroep” (appeal), “vordering” (demand) and “vonnis” (judgement) have higher frequencies in the case law corpus. Single-digit numbers (1, 2, 3 etc.) are also more common in the legal corpora, as are certain “general-looking” words like “bedoeld” (intended) and “indien” (in case). There are also words, on the other hand, which are used considerably less in the legal corpora. The most striking examples of this are the personal pronouns like “ze” (they), “we” (we) and “je” (you).

#### *2.4 Status of the characteristics found*

The figures in the former section indicate that to a certain extent it is possible to *measure* the differences between legal and general texts. The question is then how this knowledge can be put to use. One attempt to do this is described in the thesis mentioned above.<sup>5</sup> It contains a chapter in which a *similarity-measure* for documents is introduced. This measure can be calculated from certain word use statistics alone. Figures about the distribution of word types over the documents provide the essential data for this purpose. In a certain document pair, we can count the number of “hits” (word types which are used in both documents, or which are present in other parts of the corpus but are absent in both documents) and “misses” (word types which are only used in one of the two documents). Every hit and miss can be assigned a weight, based on the distribution of the corresponding word type in the corpus. This means that the presence of rare types or the absence of common types in both documents can increase their similarity considerably, whereas the presence of common types or the absence of rare types usually only has limited influence.<sup>6</sup>

Using a similarity-measure to indicate how much certain entities resemble each other is not new. The principles which underlie a similarity relation are quite general, and can be applied in many fields. In the field of linguistics, several different types of similarity have been studied in the past. An example of this can be found in Stiles (1961), where an “association factor” is defined which can be used to express the relationships between index words in a document. There are even examples of research projects where a similarity measure is used to improve the search capabilities of document retrieval systems, such as the “SMART”-project.<sup>7</sup> In this last project similarity was again only calculated for a limited number of (index) words taken from the documents in the database, however. What is special about our approach is that similarity is calculated from *all the words* which are present in a pair of documents, and even from the words which are absent from them but are present in other parts of the database. The obvious advantage of this is, of course, that no document pre-processing (in the form of the selection of index words) is necessary. Instead, it follows from the similarity algorithm (in which every word has a weight which inversely corresponds to its frequency in the corpus) which words are important when comparing certain documents and which are not. Using this algorithm, similarity can be calculated between all different pairs of documents in a corpus. The result of this is a list of all document pairs, ordered according to the measure of similarity in word use.

The final step we have taken is that we have used the similarity figures of the documents in the corpus in an experiment to form *clusters of documents* from the database automatically. The idea was that documents with a high level of similarity can be joined, to form one (larger) new document. After this step, it is necessary to calculate similarity

---

<sup>5</sup> Van Noortwijk (1995), p. 221-265.

<sup>6</sup> For an overview of actual formula which can be used to calculate similarity between two entities see Batagelj and Bren (1993).

<sup>7</sup> Salton (1971), p. 223 e.v. An overview is also given in Salton (1989).

measures again (because some documents have changed), which can again result in documents which have mutually high similarities, and therefore can be joined. The process can be repeated again and again, until at a certain moment no high similarities are found anymore (or, in a less ideal situation, some of the clusters have become unmanageably large). In the clustering process, several different techniques and criteria can be used, of which only a few were tried out by us. In general, the method has yielded promising results when applied to the statute law corpus. In this corpus a number of interesting clusters were identified, many of which contained documents which not only shared certain form characteristics, but were also related at the semantic level.

In the next section, we will describe a further possibility to create a practical computer application based on word use statistics. It can be seen as a logical next step after studying the similarity of documents. The system which is being introduced has been developed in the past few years at Erasmus University. It makes use of the characteristics of the word use in documents in order to implement a form of *conceptual legal document retrieval*.

### **3 Defining “concepts” based on word use**

#### *3.1 Introduction*

Many recent publications on legal information retrieval agree that traditional automated systems for this purpose do not satisfy the demands of lawyers. A multitude of suggestions for improvement have been made. Many of these suggestions have in common that legal information retrieval systems should be “conceptual”. Roughly speaking this means that legal information retrieval systems should contain more knowledge about the law and be more “intelligent”. In Wildemast and De Mulder (1992) an overview was given of the attempts that have been made to build such systems.

#### *3.2 Conclusions from recent attempts to implement conceptual legal information retrieval systems*

The methods proposed in literature for legal conceptual retrieval are aimed at:

- the interface with the users
- the representation of documents
- the search operation (Wildemast and De Mulder 1992).

It is the *interface* which makes communication between the user and the computer possible. It assists in the translation of the user's question into an actual search instruction for the computer. When the search instruction has been carried out, the interface is responsible for the reproduction of the results. On the basis of these results it is then possible to assess the relevance of the documents which have been found and to reformulate the question (or the actual command) if necessary. An interface can also assist the user in the formulation of the question. See for example Vries *et al.* (1991).

Conceptual retrieval can be realised by assisting the user in (re)formulating a search request. This is done by assisting the user in finding the right words to describe the concept and by providing the legal context in which concepts are described.

The *representation methods* are based on the assumption that conceptual retrieval can be realised if the representations of the original texts are based on the legal importance or legal meaning of a text.

By *search operation* is understood the function which ensures that the concrete search instruction (whether or not already re-worked in the interface) is carried out on the documents represented in the system. Most search operations (for instance the well known Boolean search) make use of the occurrence of a term rather than, for example, the term frequency in a document. The result of the Boolean search operation is the answering

of a yes/no question for each document as to whether the document satisfies the search instruction. Other search operations look for a standard which indicates the extent to which the document satisfies the search instruction. This may possibly be expressed in the form of an estimation of probability (Salton 1989), (Bookstein and Klein 1990). A similar result is achieved by search techniques which make use of "neural networks". Conceptual retrieval with the help of neural networks was proposed in Belew (1987) and Rose and Belew (1989).

The analysis of the advantages and disadvantages of the techniques presented in current literature leads to the conclusion that in most cases the method of text representation or the interface do not allow the user to define his own concepts. It would, however, be desirable that this would be the case and that these concepts could then be more precisely re-defined on the basis of the results of search operations or interpretations by the interface. Such a system could store the user's concepts: it would become a "learning" system.

As regards the interface, it is especially important that the user can bring into the system and modify his own concepts. We would argue that the quality of the interface is, therefore, the constraining factor in conceptual legal information retrieval at present. Research efforts should be concentrated on this area as a lot more can be done. For example, in the available literature hardly any attention is paid to an obvious method of allowing the user to make his own ideas explicit: the user can give the system *examples* of clearly relevant documents with which he is familiar. See Bookstein and Klein (1990).

The choice of search technique is not a crucial design decision as, given the design choices for interface and document representation, various search techniques can be used as alternatives or supplements to each other.

A prototype of a system, containing a very large collection of legal cases and formal legislation and operating with techniques in which the considerations formulated here have been realised, has been constructed at the Centre for Computers and Law of the Erasmus University, Rotterdam. We like to refer to it as a learning **concept processor**. Via the interface the documents can be looked up and given a relevance score through statistical techniques. This score helps the user to indicate which documents are relevant and which are not. Concepts are - roughly speaking - stored in terms of sets of relevant documents, with concept names, user name and date and time. Relationships between concepts can be traced and/or indicated by the users. In our opinion, such a concept processor is a necessary part of a legal conceptual retrieval system because in law concepts do not have a fixed and objective content, but can vary from user to user, from problem to problem and from time to time.

### *3.3 The system*

In principle a user of our concept processor is required to define his (legal) concepts by entering a list of documents in the database that he considers to be relevant. (These documents, identified by the user as relevant to his concept, are called "**exemplars**".) Consequently, the searching facility of the system will search for documents that are similar to the exemplars. In order to fulfil this task, the program will compare the properties or **attributes** of potentially relevant documents with those of the exemplars. These attributes consist of the words used in the documents, their frequency, possibly the order in which the words appear etc.

In section 2 we explained the possibility of measuring similarity. In a way, this principle is applied here again. By calculating the similarity between each document and all the exemplars the system is capable of ordering documents according to their relevance. Those documents that are ranked at the top of the list are the ones that the user will be interested in the most. If the system comes up with a document that the user identifies as relevant, he/she can decide to add it to the list of exemplars. The next search operation will then be based on more information than the initial one.



There is also a very important use for the documents that the system initially ranks highly, but that the user identifies as non-relevant. Although possibly for the initial stages of learning the system could compare the set of exemplars to all other documents in the total set, it is necessary that the user can provide the system with a set of “**counter exemplars**” that are as similar as possible to the exemplars of *relevant* documents, but are *irrelevant* in the given search operation. Typically, the user would inform the system that documents that are put forward as candidates for relevant documents are in fact counter exemplars. These non-relevant documents will “teach” the system the finesses of the concept the user has in mind. A concept, as used in such a conceptual retrieval system, could therefore be defined as follows:

*A concept is an ordered pair of sets of documents. The first set of the pair is the set of exemplars (of relevant documents). The second set of the pair is the set of counter exemplars (a set of non-relevant documents that are as similar as possible to the relevant documents).*

A concept can be referred to by a term that indicates the membership of the first set of the pair. For example: “(documents which contain) civil law (cases)”. The exemplars and counter exemplars of a concept together form a concept of a hierarchically “higher” order. An example would be a set of documents dealing with court cases using eye-witnesses in evidence in criminal law trials, and the set of counter exemplars consisting of documents about evidence in criminal law, but not about eye-witnesses. The hierarchically higher category would be “evidence in criminal law”. The set of exemplars would consist of both exemplars and counter exemplars of the lower concept, and the set of counter exemplars would consist of documents dealing with cases about procedural criminal law, but not about evidence (see figure 1).

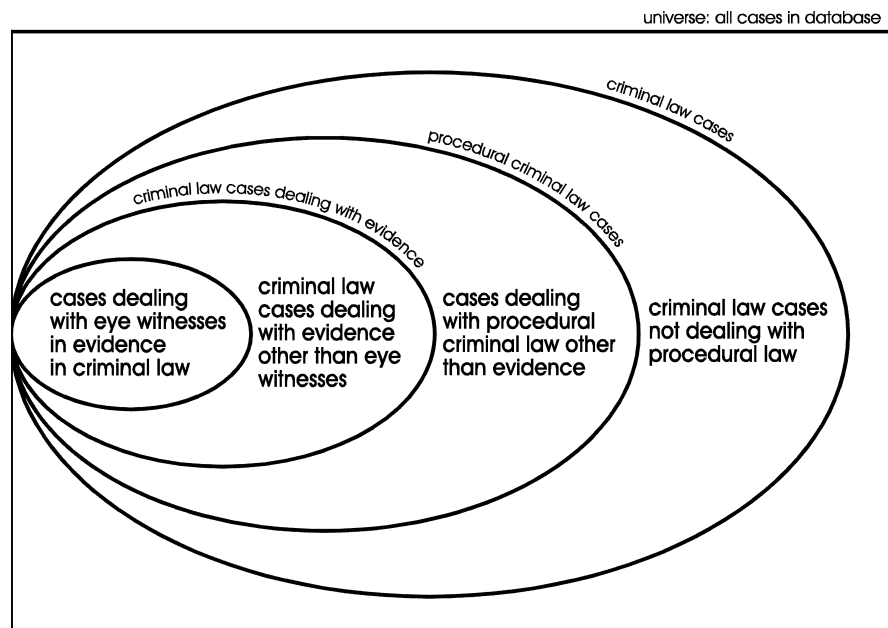


Figure 1: Hierarchy of concepts

As the “concept processor” is meant to be part of the user interface of a computer program and, furthermore, the amount of information to be processed is usually large it is important that a graphical representation of the data can be given. The use of a concept of “concept” taken from the set theory, enables such a visualisation. The graph in figure 1 shows a so called Venn-diagram which provides a clear picture of the hierarchy of concepts. The broader concepts, i.e. the ones “higher” in the hierarchy, are shown as larger ellipses.

### *3.4 Relevant documents*

Based on the model presented in the former sections, a prototype application has been created at Erasmus University which is capable of ranking documents according to their relevance to a certain concept, by comparing the word use in these documents. This system works in the following way. Suppose that a user is looking for documents about “eye witnesses in criminal law”. Within a previously selected set of documents (for instance: a certain part of a case law databank) the user has indicated a number of exemplars and counter exemplars. Based on similarities and dissimilarities in the word use of the (counter)exemplars and the other documents the computer will then calculate a measure of the probability that each of the documents is relevant. The list of documents will be ordered according to the measure of probability, and shown on the computer screen.

In an ideal situation this would immediately give a correct ranking of all documents according to the intended concept. In most cases, however, the ranking has some shortcomings. It is possible, for instance, that some of the exemplar documents are not at the top of the list, or that documents are listed amidst exemplars which on inspection are not relevant to the search concept. Most of the time it is therefore necessary to inspect some of the remaining documents, and find extra exemplars and/or counter exemplars. After some have been found, the computer can be instructed to calculate probability measures again. This will result in a new ordering of the documents, which again has to be validated by the user. Usually, the new (counter)exemplars contain a certain amount of new information about the search concept, for instance in the form of words which were not present in the earlier exemplars.

Another possibility would be that the user finds that a whole *category* of documents, which he intended to select from the database, is missing in the top part of the ranking. It is then necessary to look for at least one or two exemplars of this category, as it is probably not “covered” by the concept yet. Finally, the user will find that adding new (counter)exemplars hardly changes the ranking anymore. At this point, all relevant documents should be positioned in the top part of the ranking. The only thing left to do for the user is then to find the exact position where to “draw the line” between relevant and irrelevant documents. We have found that in some cases a graphical representation of the final probability measures can be helpful in locating this point.

## **4 Conclusion**

Studying the word use in legal documents provides empirical knowledge about the contents and the structure of legal documents. A number of techniques, many of which find their roots in the field of quantitative linguistics, can be used in this work. Results from this research can serve to develop more powerful legal information systems.

One example is a system which is capable of forming document clusters out of a legal database. This system is described in detail in Van Noortwijk (1995). The document clusters found through this technique could, among other things, play a role in increasing the recall of an information retrieval system by adding to the result of a query those documents which are in the same cluster as (and therefore probably related to) the documents which were originally found.

As a second example, a usable computer application based on word use statistics has been presented, in the form of a conceptual legal information retrieval system. This system

has been designed to give the user the possibility to define and implement his own concepts. The system provides objective feedback on the clarity and, therefore, effectiveness of the definitions given. An interesting consequence of this is that by using the system, the user not only enables the system to “learn”, but he also clarifies his own concepts and, therefore, increases his knowledge as well as learning from the end result of his search activities.

## References

- Baayen, R.H. (1989). *A corpus-based approach to morphological productivity*. Proefschrift, Centrum voor Wiskunde en Informatica (Vrije Universiteit), Amsterdam 1989.
- Bailey, R.W. (1969). “Statistics and Style: a historical survey”. In: L. Dolezel and R.W. Bailey (eds.), *Statistics and Style*, p. 217-236. Elsevier, New York 1969.
- Batagelj, V. and M. Bren (1993). *Comparing Similarity Measures*. University of Ljubljana, Ljubljana 1993.
- Belew, K. (1987). “A connectionist approach to conceptual information retrieval”. In: *Proceedings of the First International Conference on AI and Law*, Boston, 1987, pp.116-126.
- Bookstein, A. and S.T. Klein (1990). “Information Retrieval Tools for Literary Analysis”. In: Tjoa, A.M., and R. Wagner (eds.), *Database and Expert Systems Applications (DEXA)*, Proceedings of the International conference in Vienna, Austria, 1990, p. 1-7.
- Guiraud, P. (1959). *Problèmes et Méthodes de la Statistique Linguistique*. Reidel publishing company, Dordrecht 1959.
- Herdan, G. (1966). *The advanced theory of language as choice and chance*. Springer-Verlag, Berlin 1966.
- Kerkmeester, H.O. (1989). *Recht en Speltheorie (Law and Game Theory)*. With a summary in English. Vermande, Lelystad 1989.
- Kucera, H. and N.W. Francis (1967). *Computational analysis of present-day American English*. Brown University Press, Providence 1967.
- Loevinger, L. (1949). “Jurimetrics, the next step forward”. *Minn. Law Review*, april 1949, p. 455.
- Mulder, R.V. De (1984). *Een model voor juridische informatica (A Model for the application of computer science to law)*. With a summary in English. Vermande, Lelystad 1984.
- Mulder, R.V. De, M.J. van den Hoven and C. Wildemast (1993). The concept of concept in 'conceptual legal information retrieval'. In: *8th BILETA Conference Pre-proceedings*, p. 79-92. CTI Law Technology Centre, University of Warwick, Warwick 1993.
- Noortwijk, C. van (1995). *Het woordgebruik meester. Een vergelijking van enkele kwantitatieve aspecten van het woordgebruik in juridische en algemeen Nederlandse teksten (Legal word use, a comparison of some quantitative aspects of the word use in legal and general Dutch texts)*. With a summary in English. Koninklijke Vermande, Lelystad 1995.
- Rose, D.E. and R.K. Belew (1989). “Legal Information Retrieval: a Hybrid approach”. In: *Proceedings of the Second International Conference on AI and Law*, Vancouver, 1989, pp.138-146.
- Salton, G. (ed.) (1971). *The SMART retrieval system, experiments in automatic document processing*. Prentice-Hall, Englewood Cliffs N.J. 1971.
- Salton, G. (1989). *Automatic Text Processing; The transformation, Analysis, and Retrieval of Information by Computer*, Addison-Wesley Publishing Company, Reading Massachusetts, 1989.
- Vries de, W.S., H.J. van den Herik and A.H.J. Schmidt (eds.) (1991). “Separate Modelling of User-System Cooperation”. In: Breuker, J.A., R.V. De Mulder De and J.C. Hage

*JURIX'96: Kees van Noortwijk and Richard V. De Mulder*

- (eds.), *Legal Knowledge Based Systems. Model-Based Legal Reasoning*, Jurix 1991. Koninklijke Vermande BV, Lelystad, 1991, p.28-39.
- Wildemast, C.A.M. and R.V. de Mulder (1992). "Some design Considerations for a Conceptual Legal Information Retrieval System". In: Grütters, C.A.F.M., J.A.P.J. Breuker, H.J. Van den Herik, A.H.J. Schmidt, C.N.J. de Vey Mestdagh (eds), *Legal Knowledge Based Systems: Information Technology and Law*, Jurix 1992. Koninklijke Vermande BV, Lelystad 1992, pp.81-92.