Formal Concept Analysis in Information Science[*]

Uta Priss

1 Introduction

Formal Concept Analysis (FCA) is a method for data analysis, knowledge representation and information management that is widely unknown among information scientists in the USA even though this technology has a significant potential for applications. FCA was invented by Rudolf Wille in the early 80s (Wille, 1982). For the first 10 years, FCA was developed mainly by a small group of researchers and Wille's students in Germany. Because of the mathematical nature of most of the publications of that time, knowledge of FCA remained restricted to a group of "insiders". Through funded research projects, FCA was implemented in several larger-scale applications, most notably an implementation of a knowledge exploration system for civil engineering in cooperation with the Ministry for Civil Engineering of North-Rhine Westfalia (cf. Eschenfelder et al. (2000)). But these applications were not publicized widely beyond Germany.

During the last 10 years, however, FCA has grown into an international research community with applications in many disciplines, such as linguistics, software engineering, psychology, AI and information retrieval. This shift is due to a variety of factors (cf. Stumme (2002)). A few influential papers stirred interest for FCA in several fields. For example, Freeman and White's (1993) paper on social network analysis initiated an interest for the use of FCA software among sociologists. In software engineering, several FCA papers (such as Fischer (1998) and Eisenbarth et al. (2001)) won Best Paper Awards at conferences (Snelting, to appear) because FCA happened to facilitate a type of analysis which was previously not available in that field. As Stumme (2002) explains, FCA shifted emphasis to applications in computer science partly due to a merger with the Conceptual Graphs community (Sowa, 1984). An overview of the relationship between Conceptual Graphs and FCA is provided by Mineau et al. (1999).

Some of the structures of FCA appear to be fundamental to information representation and were independently discovered by different researchers. For example, Godin et al.'s (1989) use of concept lattices (which they call "Galois lattices") in information retrieval is based on an independent discovery by

[*] This is a draft version of a paper to be pulished in Cronin, Blaise (Ed.), Annual Review of Information Science and Technology, ASIST, Vol. 40. (For references to this paper, please, quote the final version, not this draft.)

Barbut & Monjardet (1970). Apart from starting a following in France and French-Canada, Godin's work also influenced Carpineto and Romano's (1993) work in Italy. Over time these groups have grown into a joint international FCA community, geographically stretching as far as Peter Eklund's group in Australia. In Europe, FCA groups can be found in Germany, France and Eastern Europe. Although FCA is probably not in the mainstream in the UK, it is used in a prestigious multi-million pound, multi-university collaboration on Advanced Knowledge Technologies (Kalfoglou et al., 2004).

So far, however, FCA is fairly unknown to information scientists in the USA. To our knowledge, only one paper on FCA has so far appeared in JASIS (Carpineto & Romano, 2000). Several FCA papers were presented at ASIST conferences but did not receive a large amount of attention. Several FCA papers have been published by the ACM but more in the field of software engineering (for example, Mili et al. (1997) or Godin & Mili (1993)). There was some interest in the use of mathematical lattices for information retrieval by Soergel (1967) and others in the 1950s and 1960s (cf. Priss (2000)). But their use of lattices is different from FCA and never lead to any implemented systems.

This ARIST chapter therefore attempts to fill a gap by both introducing information scientists to FCA and providing an overview of its current research. There are many areas of FCA that we decided not to include in this paper because they seemed less relevant for information science, such as Wolff's (2004) research on Temporal Concept Analysis; Ferre's research on Logical Concept Analysis (Ferre & King, 2004); work on performance of algorithms by Kuznetsov & Obiedkov (2002) and others; research in linguistics (cf. Priss (to appear)); research in software engineering (eg. Snelting (to appear)); educational research (eg. Hara (2002)); applications in machine learning (Kuznetsov, 2004), knowledge discovery and data mining (Valtchev et al., 2004); and older research on applications of FCA in psychology (Spangenberg & Wolff, 1993).

1.1 FCA introductions, bibliographies and software

With the exception of Wolff's (1994) and Wille's (1997c) introductions to FCA, easily available, English-language introductory material has been scarce. The mathematical foundation of FCA is described by Ganter & Wille (1999a) but only a very small portion of that book discusses topics of interest to information scientists. With respect to applications, the best overview will probably be the "State of the Art" volume of the ICFCA'03 conference, which is yet to be printed but is already available in parts on the Web. This volume will contain overviews of FCA applications in information retrieval (Carpineto & Romano, to appear), software analysis (Snelting, to appear), linguistics (Priss, to appear) and many more. Carpineto & Romano's (2004a) recent book provides an overview of FCA applications in information retrieval.

An on-line bibliography of FCA can be found at www.upriss.org.uk/fca/. This bibliography does not attempt to list individual papers, which would be an impossible task, but instead contains links to bibliographies maintained by FCA research groups and to conference proceedings. Many FCA publications of recent years appear in the proceedings of two series of international conferences: the International Conference on Conceptual Structures (ICCS), which contains FCA papers since 1995, and the International Conference on FCA (ICFCA), which was started in 2003. Many of the older publications, however, are more difficult to obtain. For example, the conference proceedings of conferences held in Germany before 2000 are mostly out of print.

An overview of FCA software is provided by Tilley (2004). A few years ago, there was no open-source FCA software available. Free software could only be obtained if one knew someone who was developing such software. Furthermore the documentation and interfaces were usually written only in German or only comprehensible for mathematicians. The availability of several Java-based open-source tools since 2003, such as ConExp (sourceforge.net/projects/conexp) and ToscanaJ (Becker & Hereth Correia, to appear) is probably another contributing factor for the recent growth of interest in FCA. These tools are cross-platform compatible, easy to install and fairly easy to use.

Current FCA software is still far from realizing the full potential of applications as described in theoretical research papers. Eklund et al. (2000) describe a possible implementation of some of the more advanced FCA aspects but to our knowledge none of that has been implemented so far. Due to the complexity of the underlying lattice data structures and of the visualizations, FCA software can usually not be developed in the kinds of short-term projects which are normally funded by national research agencies. With a few exceptions, such as the open-source software mentioned above and Carpineto & Romano's (2004b) Credo engine, which is accessible at credo.fub.it, FCA software often exists more virtually than actually.

There are at least two companies who apply and develop FCA software. The Germany-based company Navicon (www.navicon.de) was founded 10 years ago by some of Wille's former students. Their vision was to employ FCA software for information management tasks. The company is doing well but is nowadays more focused on database technologies than on FCA (F. Vogt, personal communication, February, 2003). A second commercial venture is an Australian company that is marketing an email analysis tool based on FCA (www.mail-sleuth.com).

2 Basic notions of Formal Concept Analysis

2.1 The duality of extension and intension

The basic notions of FCA are described in this section but without their mathematical details, which can be found in Ganter & Wille (1999a). A central notion of FCA is a duality called a "Galois connection". This duality can often be observed between two types of items that relate to each other in an application, such as objects and attributes or documents and terms. A Galois connection implies that if one makes the sets of one type larger, they correspond to smaller sets of the other type, and vice versa. For example, considering documents and

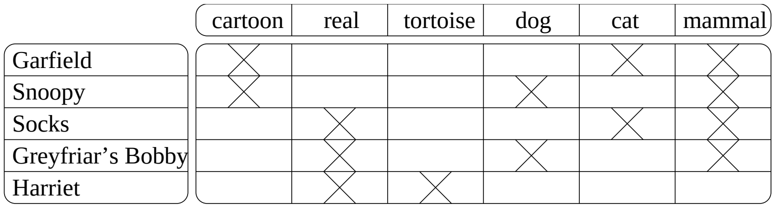

	cartoon	real	tortoise	dog	cat	mammal
Garfield	×				×	×
Snoopy	×			×		×
Socks		×			×	×
Greyfriar's Bobby		×		×		×
Harriet		×	×			

Figure 1: A formal context of "famous animals"

terms in information retrieval, enlarging a set of terms will reduce the set of documents which contain all of these terms, whereas a smaller set of terms will match a larger set of documents.

In Formal Concept Analysis, the elements of one type are called "formal objects", the elements of the other type are called "formal attributes". The adjective "formal" is used to emphasize that these are formal notions. "Formal objects" need not be "objects" in any kind of common sense meaning of "object". But the use of "object" and "attribute" is indicative because in many applications it may be useful to choose object-like items as formal objects and to choose their features or characteristics as formal attributes. In an information retrieval application, documents could be considered object-like and terms considered attribute-like. Other examples of sets of formal objects and formal attributes are tokens and types, values and data types, data-driven facts and theories, words and meanings and so on.

The sets of formal objects and formal attributes together with their relation to each other form a "formal context", which can be represented by a cross table (see figure 1). The elements on the left side are formal objects; the elements at the top are formal attributes; and the relation between them is represented by the crosses. In this example, the formal objects are animals who are famous in certain parts of the world: the cartoon characters Garfield and Snoopy, Bill Clinton's cat Socks, Greyfriar's Bobby (a dog who lived in Scotland and who is featured in British children's stories for his dedication to his deceased owner), and Harriet the tortoise (claimed to be the planet's oldest animal, brought to England by Darwin and now living in the Australia Zoo). The attributes describe whether these animals are cartoon characters or "real" animals, whether they are dogs, cats, mammals or tortoises. This is, of course, a toy example but it is sufficient to explain the basic features of FCA.

A further interesting feature of Galois connections is that a certain "closure" of the relation is implied. Starting with any set of formal objects one can identify all formal attributes which they have in common. Using the example from figure 1, the only common attribute of Harriet and Bobby is that they are both real. But they are not the only real animals, because Socks is also real. If one starts with Harriet and Bobby, derives all their shared attributes (i.e., "real") and then derives all other animals which are also real, one obtains the object set "Harriet, Bobby, Socks" and the attribute set "real". At this point the relation is "closed"

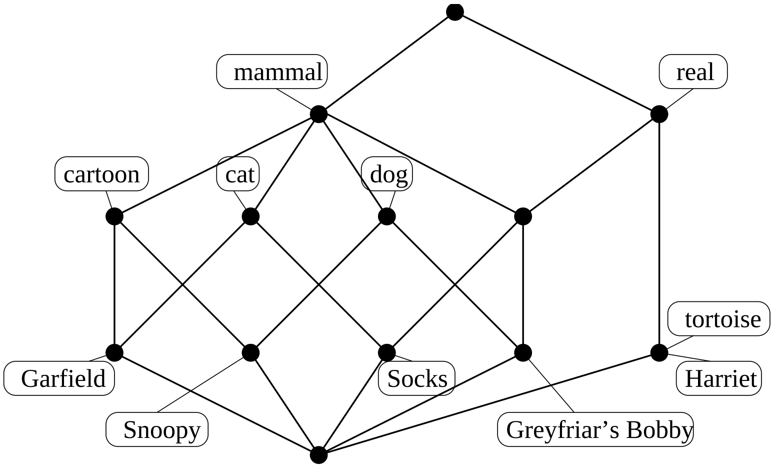

Figure 2: A concept lattice for the formal context in figure 1

because one can neither enlarge the attribute nor the object set. If one starts with Bobby and Socks, one obtains "mammal" and "real" and then no further objects or attributes because no other formal objects have both these formal attributes. A pair of a set of formal objects and a set of formal attributes that is closed in this manner is called a "formal concept". Thus "Harriet, Bobby, Socks" and "real" form a formal concept and "Bobby, Socks" and "mammal, real" form a different formal concept. The set of formal objects of a formal concept is called its "extension"; the set of formal attributes its "intension". For a given formal context, the formal concepts, their extensions and intensions are uniquely defined and fixed.

2.2 Concept lattices

An important advantage of FCA is that the Galois connections and the sets of formal concepts can be visualized. Figure 2 shows a so-called line diagram of a concept lattice corresponding to the formal context in figure 1. A concept lattice consists of the set of concepts of a formal context and the subconcept-superconcept relation between the concepts (cf. Ganter & Wille (1999a) for the mathematical details). The nodes in figure 2 represent formal concepts. Formal objects are noted slightly below and formal attributes slightly above the nodes which they label. Continuing with the example from the previous section, the node on the right side which is labeled with the formal attribute "real" shall be referred to as Concept A.

To retrieve the extension of a formal concept one needs to trace all paths which lead down from the node to collect the formal objects. In this example, the formal objects of Concept A are Socks, Bobby and Harriet. To retrieve the intension of a formal concept one needs to trace all paths which lead up in

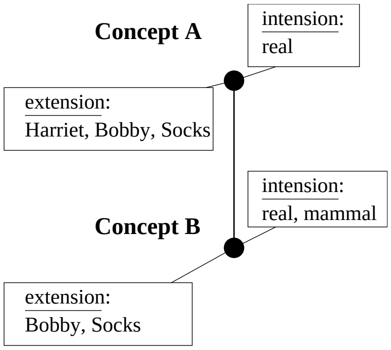

Figure 3: A subconcept-superconcept relation

order to collect all the formal attributes. In this example, there is a node above Concept A but that node has no formal attributes attached. Thus Concept A represents the formal concept with the extension "Harriet, Bobby, Socks" and the intension as the single-element set "real". The other formal concept mentioned in the last section was the one with extension "Socks, Bobby" and intension "real, mammal". This concept is connected to Concept A by an edge (the line going down from Concept A to the left) and not labeled by any object or attribute in the line diagram in figure 2. This concept shall be referred to as Concept B.

Figure 3 summarizes the relationship between Concept A and Concept B. Concept B is a subconcept of Concept A because the extension of Concept B is a subset of the extension of Concept A and the intension of Concept B is a superset of the intension of Concept A. All edges in the line diagram of a concept lattice represent this subconcept-superconcept relation.

The top and bottom concepts in a concept lattice are special. The top concept has all formal objects in its extension. Its intension is often empty but does not need to be empty. In the example in figure 2, the top concept could have a formal attribute "animal". The bottom concept has all formal attributes in its intension. If any of the formal attributes mutually exclude each other (such as "dog" and "cat") then the extension of the bottom concept must be empty (because no formal object can be a dog and cat at the same time). The top concept can be thought of as representing the "universal" concept and the bottom concept the "null" or "contradictory" concept of a formal context.

The subconcept-superconcept relation is transitive, which means that a concept is subconcept of any concept which can be reached by traveling upwards from it. If a formal concept has a formal attribute then its attributes are inherited by all its subconcepts. This corresponds to the notion of "inheritance" used in the class libraries of object-oriented modeling. For this reason, FCA is suitable for modeling and analyzing object-oriented class libraries, which has been described, for example, by Godin & Mili (1993) and by Huchard & Leblanc

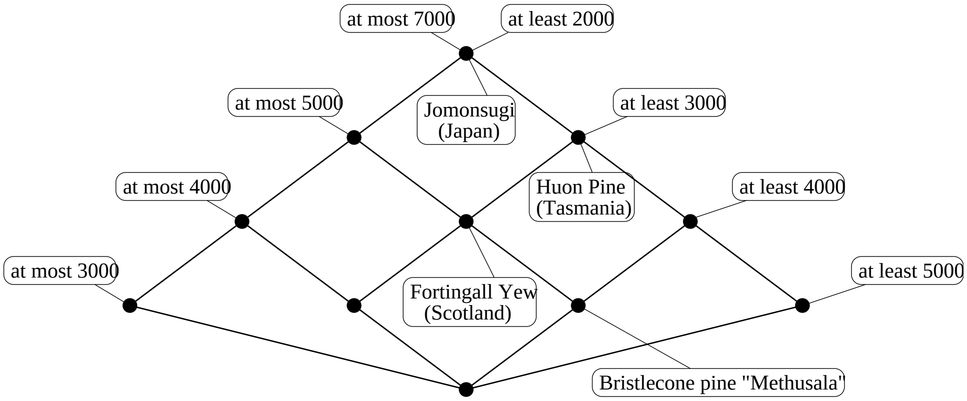

Figure 4: An interval scale for the world's oldest trees

(2000).

2.3 Conceptual scaling

A single concept lattice of large sets of objects and attributes can become fairly large and complex. Visualizations of concept lattices are only of interest if they are not too "messy" to be comprehensible for a human user. The only information which a messy visualization provides is that the underlying lattice is complex. Lindig & Snelting (1997) use this as an advantage in one case: by showing that the concept lattice of dependencies between different pieces of software code is extremely messy, they provide an argument for not attempting to re-engineer such code. But otherwise concept lattices should either be human-readable, or the information should be displayed in a different format. For example, in some applications, such as Carpineto & Romano's (2004b) Credo engine, concept lattices are mainly used as internal structures which are not displayed as a whole to the user. In Credo, only parts of the lattices are displayed similar to file/folder displays, where a second level of the hierarchy is indented and can be expanded or collapsed interactively by users. Other applications use fish-eye displays, such as described by Carpineto & Romano (1995).

A different approach for reducing the complexity of concept lattices is the idea of dividing lattices into different components based on groupings of related formal attributes (similar to the notion of "facets" in library science). These groups are then separately visualized as lattices. For some of such groups, it is possible to predict the structure of the formal attributes without even considering the formal objects. For example, the structure of data derived from a survey often represents rank orders ("agree strongly", "agree", "neutral", "disagree", "disagree strongly"). The lattice for such a rank order can be drawn without considering the actual results from the survey. Ganter & Wille (1989) call such lattices "conceptual scales". After drawing a conceptual scale, formal objects

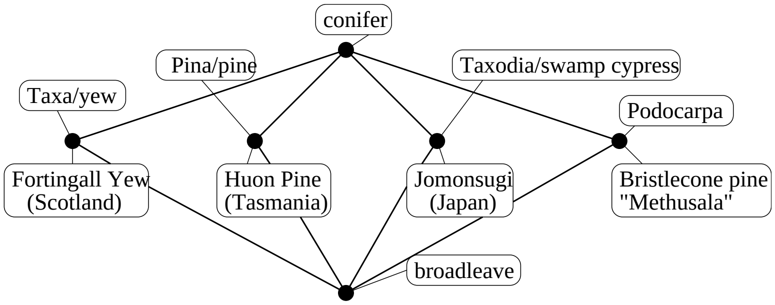

Figure 5: A nominal scale for the world's oldest trees

can be mapped to their positions on the scale. In that manner a single scale can be re-used for different formal contexts. In the survey example, the same scale could be used for different surveys to compare their results. According to Ganter & Wille (1989), typical conceptual scales are "nominal scales" (cf. figure 5), "ordinal scales" and "interval scales" (cf. figure 4). These names of these scales are adopted from traditional statistics. In addition to facilitating re-use, conceptual scales can also be utilized for analyzing dependencies between attributes (Ganter & Wille, 1989).

An interval scale as shown in figure 4 is suitable for formal attributes which have a range of possible values. In this example, the formal objects are trees which are claimed to be the world's oldest trees according to sites on the WWW. The formal attributes describe the claimed ages. For example, some sites claim that the Jomonsugi tree is 2000 year old whereas others claim that it is 7000 years old. Therefore this tree is assigned the formal attributes "at least 2000" and "at most 7000". The age range for the North American Methusala is much smaller. It is usually claimed to be about 4700 years old. Since 4000 is larger than 2000, this tree is also assigned the formal attributes "at least 2000" and "at least 3000" and so on, but its more precise values are "at least 4000" and "at most 5000". The further down a formal object is in an interval scale, the more precise its range of values is.

Figure 5 shows a second scale for the same set of formal objects as in figure 4. In this case the formal attributes are the botanical classifications of the trees. All four trees are conifers, but each belongs to a different botanical family. As mentioned above, this kind of scale, which is really just a partition of the data into separate classes, is called a "nominal scale". Several scales which are applied to the same set of objects can be combined in so-called nested line diagrams (Ganter & Wille, 1999a). The nested line diagram in figure 6 shows the scales from figures 5 and 4 combined (although in a somewhat simplified manner). The interval scale is used as inner scale, whereas the nominal scale is used as outer scale.

The purpose of the Toscana software (Kollewe et al, 1994) is to facilitate interactive exploration of such nested line diagrams. In Toscana, a user can select

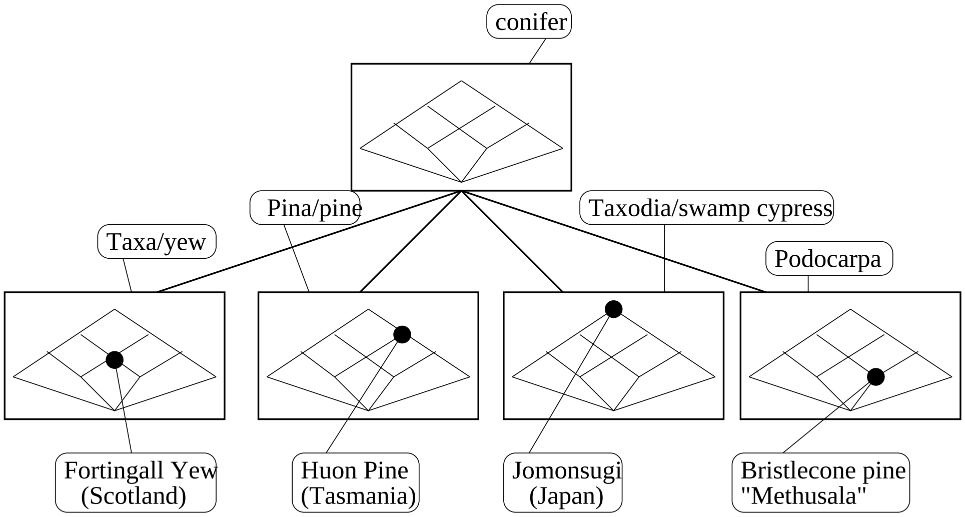

Figure 6: A nested line diagram combining figures 4 and 5

any scales that have been prepared and combine them. The user can choose the order in which the scales are combined, i.e., which are outer or inner scales. It is possible to zoom in and out of scales and to choose different parameters for the display of the formal objects. The initial set up of a "Toscana system" for a new application is quite labor-intensive because usually only some of the formal attributes can be depicted with standard scales. For qualitative data, new scales may need to be designed for every application. The Toscana software has been used in many applications so far, including some large scale applications which required the manual design of many scales. Examples are the Toscana system for the civil engineering building codes of the North-Rhine Westfalian Ministry (Eschenfelder et al. 2000) mentioned before and the Toscana system for a library catalog (Rock & Wille, 2000).

Toscana systems do not represent the only possible manner for combining scales. A different approach is suggested in the Faceted Information Retrieval (FaIR) system (Priss, 2000), where scales are not nested but presented side by side. The relationship between the scales is explored by selecting elements in one scale which will cause the corresponding elements in other scales to be highlighted.

2.4 Implications

It is beyond the scope of this paper to discuss all FCA techniques that have been used or described for applications. But an important area that must be mentioned is the use of "implications". The example in figure 2 shows some dependencies among formal attributes. For example, the formal attribute "mammal" is attached to a node that is above the nodes for "dog" and "cat". That means that the formal attributes "dog" or "cat" imply the formal attribute "mammal". This is, of course, a biological truth. But there are more attribute implications in this example some of which may hold only for the example but not in general, such as "cartoon" implies "mammal" and "tortoise" implies "real". Implications can also involve combinations of attributes. In figure 2, "cat" and "dog" together imply all other formal attributes because they are contradictory and the paths down from "cat" and "dog" meet in the bottom node.

Implications have been studied since Ganter & Wille (1986). They can be used for a step-wise computer-guided construction of conceptual knowledge called "attribute exploration" (Ganter & Wille, 1999a). Support for attribute exploration is, for example, provided by the ConImp software (Burmeister, 2000) and by the ConExp program mentioned before. Starting with a preliminary formal context, the software prompts a user with a series of questions about the relationship among attributes. In the example from figure 2, the software could ask whether all cartoon animals are mammals. The user must then either agree that this is always true or must a provide a counter example of a cartoon animal that is not a mammal. This counter example is then added to the formal context. The process stops when all possibly missing implications have been checked. Attribute exploration was studied further by Stumme (1996a, 1996b) and developed into "concept exploration" (Stumme, 1997), which can be used to explore sublattices of larger data sets. But, at least to our knowledge, concept exploration is another example of advanced FCA technology which so far only exists virtually and not in publicly available software.

3 The relationship of FCA notions to similar notions in other fields

One can argue that "formal concepts" as defined in Formal Concept Analysis describe a natural feature of information representation which is as fundamental to hierarchies and object/attribute structures as set theory or relational algebra are for relational databases. The reason for this claim is that the basic FCA structures have been rediscovered over and over by different researchers and in different settings. Gerard Salton's (1968) document/term lattices are essentially concept lattices in the sense of FCA. Interestingly, however, his discussion of lattices was removed from the second edition of his book because at that time no one could envisage applications for such lattices in information retrieval. Two years later, Barbut & Monjardet (1970) describe what they call "Galois Lattices" which are also equivalent to concept lattices. Shreider from the "Russian School of Taxonomy" (cf. Gusakova & Kuznetsov, 2002) again independently discovered essentially concept lattices. Feature structure lattices as used in linguistic componential analysis are very similar to concept lattices (cf. Dyvik's (1998) work). Last but not least, Barwise & Seligman's (1997) "classifications" are also concept lattices. The main difference between all of these independent discoveries and Formal Concept Analysis is that Wille (1982) had the vision to develop concept lattices into a complex theory with many applications, whereas others stated the structural elegance of the duality and the lattices but did not know how to exploit these in applications.

Formal concepts in FCA can be seen as a mathematical formalization of what has been called the "classical theory of concepts" in psychology/philosophy, which states that a concept is formally definable via its features. This theory has been refuted by Wittgenstein, Rosch (1973) and others but as Medin (1989, p. 1476) states: "despite the overwhelming evidence against the classical view, there is something about it that is intuitively compelling". Even though from a psychological viewpoint the classical view does not accurately represent human cognition, the classical theory nevertheless dominates the design of computerized information systems because it is much easier to implement and to manage in an electronic environment. The classical view implicitly underlies many knowledge representation formalisms used in AI and in traditional information retrieval and library systems. Even if non-classical approaches are implemented (such as cluster analysis or neural networks), the resulting concepts are still sometimes represented in the classical manner.

To avoid confusion with non-classical theories or non-mathematical versions of the classical theory, concepts in FCA are called "formal concepts". Formal concepts are primarily mathematical entities. The criticism against the classical theory of concepts is not relevant as long as FCA is used in a formal domain (such as in software engineering). But if FCA is to be used in domains that are primarily concerned with human cognition, such as psychology or linguistics, the same amount of careful modeling and caution is required for FCA as is required for statistical methods in these domains. Novice users often need to be reminded that although many formal concepts in an application may correspond to intuitive notions these users might have, not all formal concepts need to do so. Linguists might argue that formal concepts are quite different from cognitive processes relating to natural languages. This is why current FCA applications in linguistics focus more on formal structures found in lexica and dictionaries than on cognitive linguistic phenomena. Formal Concept Analysis is thus not to be understood as a formal analysis of (human) concepts, but instead as a mathematical method using "formal concepts".

It should be remarked that although the notions of "extension" and "intension" are adapted from philosophy, they may be used slightly differently in that field. For example, Frege's (1892) "Sinn" (sense) and "Bedeutung" (reference) do not correspond to extension and intension as defined by FCA. According to Frege's example, "morning star" and "evening star" both have the same referent (Venus), but a different sense ("seen in the morning" versus "seen in the evening"). In FCA, an extension occurring with respect to one formal context can have only one corresponding intension, not two different intensions. To resolve this problem, Frege's example would either be modeled in FCA by defining the intension of Venus as the set of "morning star" and "evening star" or by stating that morning star and evening star do not belong to the same formal context. According to FCA, Frege's notions might be better translated as "denotation" for "Bedeutung" and "connotation" for "Sinn" (cf. Priss (1998) for a more detailed discussion). Philosophers might disagree with these definitions but, in general, mathematical formalizations as achieved by FCA always only approximate non-formal notions held in non-mathematical disciplines. The advantage of formalizations, however, is that notions are defined with absolute precision within the formal realm and that they therefore may be implementable in software.

4 FCA in information retrieval

There was some interest in the use of lattices for information retrieval by Salton (1968) and others with respect to document/term lattices and lattices of Boolean query combinations (cf. Priss (2000) for a summary of these early attempts). But none of these resulted in practical implementations and for a long time the dominant mathematical model of information retrieval was the vector space model, excluding a lattice approach. Interest for lattices was again spurred by Godin et al. (1989) who developed an information retrieval system based on document/term lattices. Godin's system was text-based without graphical representations of the lattices but that was due to the hardware limitations of that time. He discusses fish-eye and other techniques suitable for visualizations. Godin's group continued their information retrieval applications for a while. They compared information retrieval based on concept lattices to Boolean queries and to navigation in hierarchical classifications (Godin et al. 1993a) and concluded that the performance between Boolean queries and lattice navigation was similar and both better than the use of hierarchical classification. Later their interest shifted to retrieval of software components, which differs from general information retrieval because its search space is delimited by the formal nature of programming languages. Mili et al. (1997) discovered that the use of faceted classification is not advisable for software component retrieval because the cost of developing such classifications outweighs any benefits. Furthermore according to them, controlled vocabularies (as used in a faceted classification) may be too restrictive with respect to programming languages. While these results are interesting, they should probably be considered with caution because, at least to our knowledge, no one outside of Godin's group ever tried to replicate any of these experiments.

Carpineto & Romano's (1993) research was initially influenced by Godin's work but has since then been independently advanced to a high level. Their Credo engine (Carpineto & Romano, 2004b) facilitates a lattice-based metasearch of Google results. An overview of their work and FCA applications in information retrieval in general can be found in Carpineto & Romano (2004b). In that paper, Carpineto & Romano argue that FCA can serve three purposes in information retrieval: first, FCA can support query refinement. Because a document/term lattice structures a search space into clusters of related documents, lattices can be used to make suggestions for query enlargement in cases where too few documents are retrieved and for query refinement in cases where too many documents are retrieved. Second, lattices can support an integration of querying and navigation (or browsing). An initial query identifies a start node in a document/term lattice. Users can then navigate to related nodes. Further queries are then used to "prune" a document/term lattice to help users focus their search (Carpineto & Romano, 1996a). Third, a thesaurus hierarchy can be integrated with a concept lattice - an idea which was independently discussed by different researchers (eg. Carpineto & Romano (1996b), Skorsky (1997), Priss (1997)) but is probably still not sufficiently resolved.

Apart from Credo, a second FCA application that has reached professional quality is the Mail-Sleuth software (Eklund et al., 2004). This software is marketed by an Australian company and consists of a plug-in for the MS Outlook email software, which can be used to mine large email archives. The development of this software is based on earlier research on retrieval of information from semi-structured texts (Cole & Eklund (2001) and Cole & Stumme (2000)).

In general, FCA software appears to show a promise for applications in information retrieval, however, with a few limitations. Similar to latent semantic analysis, LSA (Dumais, 2004), FCA is not suited for direct manipulation of very large data sources. It is difficult to give precise upper limits because it depends on the application. It also matters whether both the object and attribute sets are large or only one of them. FCA has been applied to thousands of documents in a small library (Rock & Wille, 2000). But presumably, it would not be possible to apply FCA (or LSA for that matter) directly to the complete Google database. But either method can be applied as a secondary tool to reorganize a set of documents resulting from a Google query (as demonstrated by the Credo engine). Since both FCA and LSA employ matrices, although LSA matrices are "many-valued contexts" in FCA terminology, it would be interesting to compare both methods more closely, which, at least to our knowledge, has not yet been done. Another area where further research should be conducted pertains to the usability of FCA software. FCA technology claims to be human-centered due to its philosophical basis but only few practical usability studies exist (such as Eklund et al. (2004)). It is to be hoped that at least Credo will be extensively tested for usability.

5 FCA as a tool for knowledge representation and knowledge discovery

Formal Concept Analysis provides a contrast to some of the traditional, statistical means of data analysis and knowledge representation because of its focus on human-centered approaches. In his first paper on FCA, Wille (1982) explains that he was influenced by H. von Hentig's (1972) concerns about the status of sciences in the modern world. Hentig's idea was 'to "restructure" theoretical developments in order to integrate, rationalize origins, connections, interpretations, and applications' (Wille, 1982, p. 447). FCA was started as an attempt at restructuring mathematical lattice theory in a manner that both facilitates communication about mathematical theory to a wider non-mathematical audience and facilitates exploitation of mathematical theory for a wide range of applications. The concept lattices of FCA serve as a means for communication, exploration and discussion which complies both with Habermas's Theory of Communicative Action and Peirce's pragmatism (cf. Wille, 1997a).

Because they facilitate discussion and exploration of conceptual structures, concept lattices can be characterized as a means of external cognition in the sense of Scaife & Rogers (1996). The use of diagrams for reasoning has been formally investigated by Dau (2004). He observes that mathematicians often include diagrams in their descriptions of mathematical facts but that normally such diagrams are not permissible as arguments themselves. By formally distinguishing between a mathematical structure and its diagrammatic representation, Dau provides a framework in which diagrams can be used for formal reasoning. Thus in addition to an intuitive notion of the importance of visualizations, such as concept lattices, Dau can even formally evaluate their usefulness within a formal framework itself.

FCA has been examined with respect to principles of knowledge representation. Wille (1997a) identifies ten functions of knowledge processing (exploring, searching, recognizing, identifying, analyzing, investigating, deciding, improving, restructuring and memorizing) and investigates how these are supported by FCA. Stumme (2002) analyses FCA with respect to Davis et al.'s (1993) five principles of knowledge representation: knowledge representations as a medium of human expression, a set of ontological commitments, a surrogate, a fragmentary theory of intelligent reasoning, and a medium for pragmatically efficient computation.

Conceptual Knowledge Discovery (Hereth et al. (2000) and Stumme et al. (1998)) is mainly supported by Toscana systems (Kollewe et al., 1994) as described in section 2. In contrast to statistical software, which attempts to provide probable answers to narrow questions, Toscana systems facilitate browsing and interactive exploration of implicit and explicit structures. Because the preparation of data for input into a Toscana system is labor-intensive and requires substantial knowledge of FCA, Toscana systems are usually compiled by an FCA expert in co-operation with a domain expert. Wille (2001) argues that this is an advantage because the processes involved in creating a conceptual representation (in the sense of FCA) encourages the discovery of implicit information and facilitates the conversion of information into knowledge. Nevertheless, the effort required for setting up Toscana systems may be a reason why their use is not more wide-spread. It should be emphasized, however, that only the preparation of a Toscana system requires expertise. End-users can utilize such a system after reading a brief introduction. A side effect of the careful set-up of a Toscana system is that it can be less error prone than some statistical methods because a careful conceptual modeling prevents data misrepresentation.

We believe that there are further research opportunities for FCA in the area of knowledge discovery which have not yet been exploited. Two workshops have been held so far on the topic of "concept lattices and knowledge discovery in databases". But many of the papers at these workshops focused on algorithmic issues (such as Kuznetsov & Obiedkov (2002)) or abstract issues (such as Wille (2001)). These issues are important but it would be more interesting to see more realistic applications. Furthermore, it might be interesting to compare "relational scaling" as described by Prediger & Stumme (1999) to methods employed in business intelligence and data warehousing because they appear to pursue similar goals.

6 Applications of FCA in logic and AI

Since about 1996, attempts have been made to combine FCA with other formalisms of conceptual structures, such as Sowa's (1984) Conceptual Graphs. Wille (1997b) describes a translation of Conceptual Graphs into formal contexts and concept lattices. Mineau et al. (1999) investigate the commonalities between both theories at a more general level. Conceptual Graphs (Sowa, 1984) are a formalism for knowledge representation, which is similar to semantic networks, entity relationship diagrams and the Semantic Web standard RDF. Sowa developed his Conceptual Graphs based on Peirce's Existential Graphs, a graphical, symbolic notation for reasoning which incorporates aspects of context and modalities. In contrast to the hierarchical relations which are expressed in concept lattices, Conceptual Graphs can be used to represent semantic relations, such as part/whole, and linguistic case relations, such as Agent (or subject of a sentence), Patient (or direct object of a sentence), and so on. Natural language sentences can be translated more or less directly into Conceptual Graphs. A connection with FCA is established via types: each concept of a Conceptual Graph contains information about its type. The types form a hierarchy which can be modeled with FCA. According to Wille (1997b) a combination of FCA and Conceptual Graphs can facilitate a formalization of Elementary Logic and thus presents a powerful formalism for the representation and analysis of human reasoning and argumentation.

Apart from Conceptual Graphs, connections have been established between FCA and Description Logics (Prediger & Stumme, 1999). In contrast to Conceptual Graphs and FCA, which primarily focus on representations, Description Logics investigate expressivity and computability of logical representations. For example, Description Logics can check whether a concept subsumes another concept (with respect to a knowledge base) or whether an instance belongs to a concept.

A combination of FCA and Conceptual Graphs is not just another ontology formalism because of its philosophical foundation. Wille (2000a) perceives logic (and human reasoning and argumentation) as a Kantian triad of concepts, judgments and conclusions. The goal is to use FCA to achieve a mathematization of these three philosophical doctrines in a framework of "Contextual Logic" (Prediger, 1998). While FCA is used to mathematize concepts, Conceptual Graphs are used to mathematize judgments. A combination of Conceptual Graphs, FCA and Description Logics can then be used to mathematize conclusions. Wille sees this as a continuation of Boole's logic of signs and classes (Wille (2000b) and (2004)). But in contrast to Boole, who envisioned a "Universal" set or class, FCA focuses on formal contexts, which are finite in most applications (Ganter & Wille, 1999b) and which avoid some of the confusion caused by the assumption of universality.

A challenge for this kind of mathematization of logic is the treatment of negation (cf. Wille (2000b), Dau (2000), Kwuida et al. (2004)). For example, what is the meaning of a negated concept, such as "not a piano", or of a negated attribute, such as "not green"? Intuitively, one might in both cases construct some hypothetical superconcept and then negate with respect to that concept. "Not a piano" might be other musical instruments which are not pianos. "Not green" might be other colors. Without such a superconcept, negation would be too ambiguous. For example, one would probably not accept "library science" in the extension of a concept called "not piano".

With respect to concept lattices it can occur that negations of formal concepts are not materialized as formal concepts themselves. This is due to the fact that a formal context does not normally explicitly specify negation. If an object does not have an attribute assigned, it can mean that the attribute is irrelevant, the relationship is unknown or that the object does sometimes, usually or never have the attribute (Burmeister & Holzer, 2000). Therefore, adding negation to a concept lattice requires considering so-called protoconcepts or semi-concepts (Wille 2000b), which are mathematical structures similar to formal concepts but whose mathematical properties are more complex and more difficult to describe. The study of the structures arising from such proto- or semi-concepts is on-going (eg. Hereth Correia & Klinger, 2004). Other aspects of Contextual Logic pertain to implicit knowledge, incorporation of background knowledge (Hereth Correia & Klinger, 2004) and the incorporation of existential quantifiers (Wille, 2002).

It remains to be seen how much of the framework of Contextual Logic will eventually lead to implementations in computational ontology software or whether the focus of this research will remain on mathematical structures and philosophical implications. Some suggestions for implementations were made by Groh & Eklund (1999) and Eklund et al. (2000), but as far as we know those suggestions have not yet been instantiated. But, in any case, this research provides an interesting contrast to the commercially driven developments within the Semantic Web community.

7 Conclusion

It is hoped that this overview of formal concept analysis might stir more interest in this topic among information scientists. FCA provides a large set of methods, of which a fairly small, but powerful set has so far been implemented in software that is readily available. Applications demonstrate the usefulness of such software. The overview in this paper can only provide a glimpse of what is available for and possible with FCA.

References

Barbut M., & Monjardet B. (1970). Ordre et Classification. Algebre et Combinatoire. Tome II, Paris: Hachette.

Barwise, J., & Seligman, J. (1997). Information Flow. The Logic of Distributed Systems. Cambridge University Press.

Becker, P., & Hereth Correia, J. (to appear). The ToscanaJ Suite for implementing Conceptual Information Systems. In: Formal Concept Analysis - State of the Art, Proceedings of the First International Conference on Formal Concept Analysis. Berlin: Springer.

Burmeister, P. (2000). ConImp - Ein Programm zur Formalen Begriffsanalyse. In G. Stumme, & R. Wille (Eds.), Begriffliche Wissensverarbeitung. Methoden und Anwendungen. Berlin: Springer, 25-56.

Burmeister, P., & Holzer, R. (2000). On the Treatment of Incomplete Knowledge in Formal Concept Analysis. In B. Ganter, & G. Mineau, (Eds.), Conceptual Structures: Logical, Linguistic and Computational Issues. LNAI 1867. Berlin: Springer, 385-398.

Carpineto, C., & Romano, G. (1993). GALOIS: An order-theoretic approach to conceptual clustering. Proc. of the 10th Conf. on Mach. Learn., Amherst, MA, Kaufmann, 33-40.

Carpineto, C., & Romano, G. (1995). Ulysses: a lattice-based multiple interaction strategy retrieval interface. In B. Blumenthal, J. Gornostaev & C. Unger (Eds.), Human-Computer Interaction, 5th International Conference, EWHCI'95, Moscow, Russia, July 1995, Selected Papers, LNCS 1015. Berlin: Springer, 91-104.

Carpineto, C., & Romano, G. (1996a). Information retrieval through hybrid navigation of lattice representations. International Journal of Human-Computer Studies, 45, 5, 553-578.

Carpineto, C., & Romano, G. (1996b). A lattice conceptual clustering system and its application to browsing retrieval. Machine Learning, 24, 2, 1-28.

Carpineto, C., & Romano, G. (2000). Order-theoretical ranking. Journal of American Society for Information Science, 51, 7, 587-601.

Carpineto, C., & Romano, G. (2004a). Concept Data Analysis: Theory and Applications. John Wiley & Sons.

Carpineto, C., & Romano, G. (2004b). Exploiting the Potential of Concept Lattices for Information Retrieval with CREDO. Journal of Universal Computing, 10, 8, 985-1013.

Carpineto, C., & Romano, G. (to appear). Using concept lattices for text retrieval and mining. In: Formal Concept Analysis - State of the Art, Proceedings of the First International Conference on Formal Concept Analysis. Berlin: Springer.

Cole, R., & Stumme, G. (2000). CEM - A Conceptual Email Manager. In B. Ganter, & G. Mineau (Eds.), Conceptual Structures: Logical, Linguistic and Computational Issues. LNAI 1867. Berlin: Springer, 438-452.

Cole, R., & Eklund, P. (2001). Browsing Semi-structured Web Texts Using Formal Concept Analysis. In H. Delugach, & G. Stumme (Eds.), Conceptual

Structures: Broadening the Base. LNAI 2120. Berlin: Springer, 319-332.

Dau, F. (2000). Negations in Simple Concept Graphs. In B. Ganter, & G. Mineau (Eds.), Conceptual Structures: Logical, Linguistic and Computational Issues. LNAI 1867. Berlin: Springer, 263-276.

Dau, F. (2004). Types and Tokens for Logic with Diagrams In K. E. Wolff, H. Pfeiffer, & H. Delugach (Eds.), Conceptual Structures at Work: 12th International Conference on Conceptual Structures. Berlin: Springer, 62-93.

Davis, R., Shrobe, H., & Szolovits, P. (1993). What is a knowledge representation? AI Magazine, 14, 1, 17-33.

Dumais, S. (2004). Latent semantic analysis. Annual Review of Information Science and Technology (ARIST), 38.

Dyvik, H. (1998). A Translational Basis for Semantics. In H. Johansson, & S. Oksefjell (Eds.), Corpora and Crosslinguistic Research: Theory, Method and Case Studies, Rodopi, 51-86.

Eisenbarth, T., Koschke, R., & Simon, D. (2001). Feature-driven program understanding using concept analysis of execution trace. In Proceedings of the Ninth International Workshop on Program Comprehension. International Conference on Software Maintenance.

Eklund, P., Groh, B. Stumme, G., & Wille, R. (2000). A Contextual-Logic Extension of TOSCANA. In B. Ganter, & G. Mineau (Eds.), Conceptual Structures: Logical, Linguistic and Computational Issues. LNAI 1867. Berlin: Springer, 453-467.

Eklund, P., Ducrou, J., & Brawn, P. (2004). Concept Lattices for Information Visualization: Can Novices Read Line Diagrams? In P. Eklund (Ed.), Concept Lattices: Second International Conference on Formal Concept Analysis, LNCS 2961. Berlin: Springer, 14-27.

Eschenfelder, D., Kollewe W., Skorsky, M., & Wille, R. (2000). Ein Erkundungssystem zum Baurecht: Methoden der Entwicklung eines TOSCANA-Systems. In G. Stumme, & R. Wille (Eds.), Begriffliche Wissensverarbeitung. Methoden und Anwendungen. Berlin: Springer, 254-272.

Ferre, S., & King, R. (2004). In P. Eklund (Ed.), Concept Lattices: Second International Conference on Formal Concept Analysis, LNCS 2961. Berlin: Springer, 47-54.

Fischer, B. (1998). Specification-Based Browsing of Software Component Libraries. Proc. Automated Software Engineering, Hawaii, 246-254.

Freeman, L., & White, D. (1993). Using Galois lattices to represent network data. Sociological Methodology, 127-146.

Frege, G. (1892). Über Sinn und Bedeutung. Zeitschr. f. Phil. und phil. Kritik, 100, 25-50.

Ganter, B., & Wille, R. (1986). Implikationen und Abhangigkeiten zwischen Merkmalen. In P. O. Degens, H.-J. Hermes, & O. Opitz (Eds.), Die Klassifikation und ihr Umfeld. Frankfurt: Indeks, 171-185.

Ganter, B., & Wille, R. (1989). Conceptual scaling. In F. Roberts (Ed.), Applications of combinatorics and graph theory to the biological and social sciences. Berlin: Springer, 139-167.

Ganter, B., & Wille, R. (1999a). Formal Concept Analysis. Mathematical Foundations. Berlin: Springer.

Ganter, B., & Wille, R. (1999b). Contextual Attribute Logic. In W. Tepfenhart, & W. Cyre (Eds.), Conceptual Structures: Standards and Practices. Proceedings of the 7th International Conference on Conceptual Structures, LNAI 1640. Berlin: Springer, 377-388.

Godin, R., Gecsei, J., & Pichet, C. (1989). Design of browsing interface for information retrieval. In N. J. Belkin, & C. J. van Rijsbergen (Eds.), Proc. SIGIR '89, 32-39.

Godin, R., Missaoui, R., & April, A. (1993a). Experimental comparison of navigation in a Galois lattice with conventional information retrieval methods. Int. J. Man-Machine Studies 38, 747-767.

Godin, R., & Mili, H. (1993). Building and maintaining analysis-level class hierarchies using Galois lattices. OOPSLA '93. ACM Sigplan Notices, 28, 10, 394-410.

Groh, B., & Eklund, P. (1999). Algorithms for Creating Relational Power Context Families from Conceptual Graphs. In W. Tepfenhart, & W. Cyre (Eds.), Conceptual Structures: Standards and Practices. Proceedings of the 7th International Conference on Conceptual Structures, LNAI 1640. Berlin: Springer, 389-400.

Gusakova, S. M., & Kuznetsov, S. O. (2002). On Moscow St. Petersburg Taxonomy Schools. In G. Angelova, D. Corbett, & U. Priss (Eds.), Foundations and Applications of Conceptual Structures Contributions to ICCS 2002. Bulgarian Academy of Sciences, 92-101.

Hara, N. (2002). Analysis of Computer-Mediated Communication using Formal Concept Analysis as a Visualizing Methodology. Journal of Educational Computing Research, 26, 25-49.

V. Hentig, H. (1972). Magier oder Magister? Über die Einheit der Wissenschaft im Verständigungsprozess. Stuttgart: Klett.

Hereth, J., Stumme, G., Wille, R., & Wille, U. (2000). Conceptual Knowledge Discovery in Data Analysis. In B. Ganter, & G. Mineau (Eds.), Conceptual Structures: Logical, Linguistic and Computational Issues. LNAI 1867. Berlin: Springer, 421-437.

Hereth Correia, J., & Klinger, J. (2004). Protoconcept Graphs: The Lattice of Conceptual Contents. In P. Eklund (Ed.), Concept Lattices: Second International Conference on Formal Concept Analysis, LNCS 2961. Berlin: Springer, 14-27.

Huchard, M., & Leblanc, H. (2000). Computing interfaces in Java. ASE'00: Fifteenth IEEE International Conference on Automated Software Engineering, 317-320.

Kalfoglou, Y., Dasmahapatra, S., & Chen-Burger, Y (2004). FCA in Knowledge Technologies: Experiences and Opportunities. In P. Eklund (Ed.), Concept Lattices: Second International Conference on Formal Concept Analysis, LNCS 2961. Berlin: Springer, 252-260.

Kollewe, W., Skorsky, M., Vogt, F., & Wille, R. (1994). TOSCANA - ein Werkzeug zur begrifflichen Analyse und Erkundung von Daten. In R. Wille, &

M. Zickwolff (Eds.), Begriffliche Wissensverarbeitung - Grundfragen und Aufgaben. Mannheim: B.I.-Wissenschaftsverlag, 267-288.

Kuznetsov, S., & Obiedkov, S (2002). Comparing Performance of Algorithms for Generating Concept Lattices. Journal of Experimental & Theoretical Artificial Intelligence, 14, 189-216.

Kuznetsov, S. (2004). Machine Learning and Formal Concept Analysis. In P. Eklund (Ed.), Concept Lattices: Second International Conference on Formal Concept Analysis, LNCS 2961. Berlin: Springer, 287-312.

Kwuida, L., Tepavcevic, A., & Seselja, B. (2004). Negation in Contextual Logic. In K. E. Wolff, H. Pfeiffer, & H. Delugach (Eds.), Conceptual Structures at Work: 12th International Conference on Conceptual Structures. Berlin: Springer, 227-241.

Lindig, C., & Snelting, G. (1997). Assessing Modular Structure of Legacy Code Based on Mathematical Concept Analysis. Proceedings of the 19th international conference on Software engineering, Boston, MA, USA, 349-359.

Medin, D. L. (1989). Concepts and Conceptual Structure. American Psychologist, 44, 12, 1469-1481.

Mili, H., Ah-Ki, E., Godin, R., & Mcheick, H. (1997). Another nail to the coffin of faceted controlled-vocabulary component classification and retrieval. ACM SIGSOFT Software Engineering Notes, 22, 3, 89-98.

Mineau, G., Stumme, G., & Wille, R (1999). Conceptual Structures Represented by Conceptual Graphs and Formal Concept Analysis. In W. Tepfenhart, & W. Cyre (Eds.), Conceptual Structures: Standards and Practices. Proceedings of the 7th International Conference on Conceptual Structures, LNAI 1640. Berlin: Springer, 423-441.

Prediger, S. (1998). Kontextuelle Urteilslogik mit Begriffsgraphen. Ein Beitrag zur Restrukturierung der mathematischen Logik. PhD Thesis. Aachen: Shaker.

Prediger, S., & Stumme, G. (1999). Theory-driven Logical Scaling. Conceptual Information Systems meet Description Logics. In P. Lambrix, A. Borgida, M. Lenzerini, R. Muller, & P. Patel-Schneider (Eds.), Proceedings DL'99. CEUR Workshop Proc. 22.

Prediger, S., & Wille, R. (1999): The Lattice of Concept Graphs of a Relationally Scaled Context. In W. Tepfenhart, & W. Cyre (Eds.), Conceptual Structures: Standards and Practices. Proceedings of the 7th International Conference on Conceptual Structures, LNAI 1640. Berlin: Springer, 401-414.

Priss, U. (1997). A Graphical Interface for Document Retrieval Based on Formal Concept Analysis. In E. Santos (Ed.), Proceedings of the 8th Midwest Artificial Intelligence and Cognitive Science Conference. AAAI Technical Report CF-97-01, 66-70.

Priss, U. (1998). Relational Concept Analysis: Semantic Structures in Dictionaries and Lexical Databases, (PhD Thesis). Aachen: Shaker, 1998.

Priss, U. (2000). Lattice-based Information Retrieval. Knowledge Organization, 27, 3, 132-142.

Priss, U. (to appear). Linguistic Applications of Formal Concept Analysis. In: Formal Concept Analysis - State of the Art, Proceedings of the First

International Conference on Formal Concept Analysis. Berlin: Springer.

Rock, T., & Wille, R. (2000). Ein TOSCANA-Erkundungssystem zur Literatursuche. In G. Stumme & R. Wille (Eds.), Begriffliche Wissensverarbeitung. Methoden und Anwendungen. Berlin: Springer, 239-253.

Rosch, E. (1973). Natural categories. Cognitive Psychology, 4, 328-350.

Salton, G. (1968). Automatic Information Organization and Retrieval. New York: McGraw-Hill.

Scaife, M., & Rogers, Y. (1996). External cognition: how do graphical representations work. International Journal of Human-Computer Studies, 45, 185-213.

Skorsky, M. (1997). Graphische Darstellung eines Thesaurus. Deutscher Dokumentartag, Regensburg.

Snelting, G. (to appear). Concept Lattices in Software Analysis. In: Formal Concept Analysis - State of the Art, Proceedings of the First International Conference on Formal Concept Analysis. Berlin: Springer.

Soergel, D. (1967). Mathematical Analysis of Documentation Systems. Information Storage and Retrieval, 3, 129-173.

Sowa, J. (1984). Conceptual Structures: Information Processing in Mind and Machine. Reading: Addison-Wesley.

Spangenberg, N., & Wolff, K.E. (1993). Datenreduktion durch die Formale Begriffsanalyse von Repertory Grids. In J. W. Scheer, & A. Catina (Eds.), Einführung in die Repertory Grid Technik. Klinische Forschung und Praxis 2. Verlag Hans Huber, 38-54.

Stumme, G. (1996a). Attribute Exploration with Background Implications and Exceptions. In H.-H. Bock, & W. Polasek (Eds.), Data Analysis and Information Systems. Statistical and Conceptual approaches. Proc. GfKl'95. Studies in Classification, Data Analysis, and Knowledge Organization, 7. Berlin: Springer, 457-469.

Stumme, G. (1996b). Exploration Tools in Formal Concept Analysis. In E. Diday, Y. Lechevallier, & O. Opitz (Eds.), Ordinal and Symbolic Data Analysis. Proc. OSDA'95. Studies in Classification, Data Analysis, and Knowledge Organization 8. Berlin: Springer, 31-44.

Stumme, G. (1997). Concept Exploration - A Tool for Creating and Exploring Conceptual Hierarchies. In D. Lukose, H. Delugach, M. Keeler, L. Searle, & J. F. Sowa (Eds.), Conceptual Structures: Fulfilling Peirce's Dream. Proc. ICCS'97. LNAI 1257. Berlin: Springer, 318-331.

Stumme, G., Wille, R., & Wille U. (1998). Conceptual knowledge discovery in databases using formal concept analysis methods. In J. M. Zytkow, & M. Quafofou (Eds.), Principles of Data Mining and Knowledge Discovery. LNAI 1510. Berlin: Springer, 450-458.

Stumme, G. (2002). Formal Concept Analysis on Its Way from Mathematics to Computer Science. In U. Priss, D. Corbett, & G. Angelova (Eds.), Conceptual Structures: Integration and Interfaces, 10th International Conference on Conceptual Structures, LNCS 2393. Berlin: Springer, 2-19.

Tilley, T. (2004). Tool Support for FCA. In P. Eklund (Ed.), Concept Lattices: Second International Conference on Formal Concept Analysis, LNCS

2961. Berlin: Springer, 104-111.

Valtchev, P., Missaoui, R., & Godin, R. (2004). Formal Concept Analysis for Knowledge Discovery and Data Mining: The New Challenges. In P. Eklund (Ed.), Concept Lattices: Second International Conference on Formal Concept Analysis, LNCS 2961. Berlin: Springer, 352-371.

Wille, R. (1982). Restructuring lattice theory: an approach based on hierarchies of concepts. In I. Rival (Ed.), Ordered sets. Reidel, Dordrecht-Boston, 445-470.

Wille, R. (1992). Concept lattices and conceptual knowledge systems. Computers & Mathematics with Applications, 23, 493-515.

Wille, R. (1997a). Conceptual landscapes of knowledge: a pragmatic paradigm for knowledge processing. In W., Gaul, & Locarek-Junge (Eds.), Classification in the Information Age. Berlin: Springer, 1999, 344-356.

Wille, R. (1997b). Conceptual Graphs and Formal Concept Analysis. In D. Lukose, H. Delugach, M. Keeler, L. Searle, & J. F. Sowa (Eds.), Conceptual Structures: Fulfilling Peirce's Dream. Proc. ICCS'97. LNAI 1257. Berlin: Springer, 290-303.

Wille, R. (1997c). Introduction to formal concept analysis. In G. Negrini (Ed.), Modelli e modellizzazione. Models and modelling. Consiglio Nazionale delle Ricerche, Instituto di Studi sulli Ricerca e Documentazione Scientifica, Roma, 39-51.

Wille, R. (2000a). Contextual logic summary. In G. Stumme (Ed.), Working with Conceptual Structures. Contributions to ICCS 2000. Aachen: Shaker, 265-276.

Wille, R. (2000b). Boolean Concept Logic. In B. Ganter, & G. Mineau (Eds.), Conceptual Structures: Logical, Linguistic and Computational Issues. LNAI 1867. Berlin: Springer, 317-331.

Wille, R. (2001). Why Can Concept Lattices Support Knowledge Discovery in Databases? In E. Mephu Nguifo (Eds.), ICCS 2001 International Workshop on Concept Lattice-Based Theory, Methods and Tools for Knowledge Discovery in Databases. Stanford University, Palo Alto (CA), July 30th, 2001, 7-20.

Wille, R. (2002). Existential Concept Graphs of Power Context Families. In U. Priss, D. Corbett, & G. Angelova (Eds.), Conceptual Structures Integration and Interfaces. ICCS 2002. LNAI 2393, Berlin: Springer, 382-395.

Wille, R. (2004). Preconcept Algebras and Generalized Double Boolean Algebras. In P. Eklund (Ed.), Concept Lattices: Second International Conference on Formal Concept Analysis, LNCS 2961. Berlin: Springer, 1-13.

Wolff, K. E. (1994). A first course in formal concept analysis - how to understand line diagrams. In F. Faulbaum (Ed.), SoftStat'93, Advances in statistical software, Volume 4, Gustav Fischer Verlag, 429-438.

Wolff, K. E. (2004). Particles and Waves as Understood by Temporal Concept Analysis, In K. E. Wolff, H. Pfeiffer, & H. Delugach (Eds.), Conceptual Structures at Work: 12th International Conference on Conceptual Structures. Berlin: Springer, 126-141.