Formal Concept Analysis on its Way from Mathematics to Computer Science

Gerd Stumme

Institut für Angewandte Informatik und Formale Beschreibungsverfahren AIFB,
Universität Karlsruhe, D–76128 Karlsruhe, Germany
www.aifb.uni-karlsruhe.de/WBS/gst; stumme@aifb.uni-karlsruhe.de

Abstract. In the last years, the main orientation of Formal Concept Analysis (FCA) has turned from mathematics towards computer science. This article provides a review of this new orientation and analyzes why and how FCA and computer science attracted each other. It discusses FCA as a knowledge representation formalism using five knowledge representation principles provided by Davis, Shrobe, and Szolovits [15]. It then studies how and why mathematics-based researchers got attracted by computer science. We will argue for continuing this trend by integrating the two research areas FCA and Ontology Engineering.

1 Introduction

Formal Concept Analysis (FCA) has observed a major change of orientation in the last years. Having been introduced as a mathematization of the concept of 'concept' in the early 1980ies, its main orientation has turned from mathematics towards computer science during the last ten years: ten years ago, virtually all FCA papers were given at mathematics conferences, while nowadays they are given almost exclusively at conferences related to computer science. FCA is now considered as the mathematical backbone of *Conceptual Knowledge Processing (CKP)*, a theory located in computer science, having as task to provide methods and tools for human–oriented, concept–based knowledge processing. Seven years after the first FCA papers presented at an ICCS conference, it is time to review this trend.

In this paper, the change of orientation will be reviewed from a subjective point of view. During his stay at the Department of Mathematics at Darmstadt University of Technology and at computer science groups at Blaise Pascal University, Clermont-Ferrand, and the University of Karlsruhe, the author has observed and also actively shaped this new orientation. It will be analyzed why FCA became attractive as a knowledge representation method for computer science, and why computer science became attractive for researchers working on FCA. We start with the analysis of why FCA is a suitable knowledge representation formalism, based on the article "What is a knowledge representation?" by R. Davis, H. Shrobe, and P. Szolovits [15].

Having analyzed the attractiveness of FCA as a knowledge representation method for computer science, we will discuss why computer science became attractive for researchers working on FCA; and how FCA found a new home in computer science. The new home is *Conceptual Knowledge Processing*. Its aim is to provide methods and tools for acquiring, reasoning with, and representing knowledge, and for making it available to human communication. Currently, two main research trends can be distinguished in CKP: Contextual Logic and Conceptual Knowledge Discovery. We will discuss these two research trends, with a focus on the latter.

Section 2 provides a discussion about knowledge representation with FCA according to the principles given in [15]. In Section 3 we review the change of orientation of FCA towards computer science. Its extension to Conceptual Knowledge Processing and Discovery is the topic of Section 4. Section 5 concludes the article.

2 Knowledge Representation with Formal Concept Analysis

The convergence of FCA with computer science demands for a discussion about their relationships. In [85, 84, 69, 41, 31, 83], several aspects of this relationship have been studied. In this paper we take up the discussion. In [15], R. Davis, H. Shrobe, and P. Szolovits studied the question "What is a knowledge representation?" They provided five principles a knowledge representation should follow. We will use these principles to "characterize and make explicit the 'spirit' of [Formal Concept Analysis], the important set of ideas and inspirations that lie behind [...] the concrete machinery used to implement the representation." [15]. According to the authors, a knowledge representation is (i) a medium of human expression, (ii) a set of ontological commitments, (iii) a surrogate, (iv) a fragmentary theory of intelligent reasoning, and (v) a medium for pragmatically efficient computation.[1] The authors claim that these principles offer a framework for making explicit the 'spirit' of a representation, and the way it emphasizes on one or more of them characterizes the fundamental 'mindset' of the representation. Each knowledge representation formalism is in some way a trade-off between these principles. We will use these five criteria for discussing the role of FCA as knowledge representation method.

It will turn out that the first three principles (especially the first one) have been the driving forces for the development of FCA, while interest on the last two principles — although not completely absent at the beginning (see for instance knowledge acquisition with attribute exploration, implicational theories, and efficient computation of concept lattices [20]) — increased during the change of orientation of FCA towards computer science.

[1] Davis *et al* discuss these principles in the order 3–2–4–5–1. Here we reorder them to follow more closely the historical development of FCA.

2.1 FCA as a medium of human expression

"Knowledge representations are [...] the medium of expression and communication in which we tell the machine (and perhaps one another) about the world. [...] Knowledge representation is thus a medium of expression and communication for the use by *us*" [15]. In other words: "A representation is the language in which we communicate, hence we must be able to speak it without heroic effort".

This observation has always been predominant for the development of theory for and applications of FCA, as the strong emphasis on its philosophical roots shows. When introducing FCA in [74], R. Wille's purpose was to restructure lattice theory: "*Restructuring lattice theory* is understood as an attempt to unfold lattice-theoretical concepts, results, and methods in a continuous relationship with their surroundings [...]. One basic aim is to promote better communication between lattice theorists and potential users of lattice theory" [74, pp. 447]. The program of restructuring lattice theory followed a programmatic discussion about the role of sciences in our society by H. von Hentig [29]. Hentig requests that the sciences "uncover their non-intended aims, declare their intended aims, select and adjust their means according to those aims, discuss openly and understandably their justifications, expectations, and possible consequences, and therefore disseminate their means of research and results in common language" [29, pp. 136 f; translated by the author]. As application, Wille referred to the roots of the lattice idea, namely hierarchies of concepts, which played an important role in attempts to formalize logic [50]. Wille discusses in his visionary article "how parts of arithmetic, structure and representation theory of lattices may be developed out of problems and questions which occur within the analysis of contexts and their hierarchies of concepts" [74, pp. 448].

A second philosphical foundation of FCA is the pragmatic philosophy of Ch. S. Peirce [42], and the Theory of Communicative Action of J. Habermas [26] (cf. [78, 81]). Peirce considers knowledge as always incomplete, formed and continuously assured by human discourse. J. Habermas took up these ideas in his Theory of Communicative Action where he emphasizes on the importance of the inter-subjective community of communication. He observes that humans operate in argumentative dispute on the normative basis of practical-ethical rules. Even in scientific statements (i. e., in assertions), one tries to convince the listener and expects agreement or counter-arguments. Hence even in these apparently objective domains the ethical norms of equality and acceptance are thus present (cf. [32, p. 338]). Following this line of argumentation, the task for theories formalizing aspects of knowledge is thus to provide means for rational communication. The observation that this understanding conflicts with the widely accepted view of mathematics as a means for mechanistic problem solving was certainly one of the main reasons for the change of orientation of FCA towards computer science, where human(–computer) interaction is considered as a research topic on its own (although large parts of computer science also follow a rather mechanistic view).

2.2 The ontological commitment of FCA

Knowledge Representation "is a *set of ontological commitments*, i. e., an answer to the following question: In what terms should I think about the world? [...] In selecting any representation, we are [...] making a set of decisions about how and what to see in the world. [...] We (and our reasoning machines) need guidance in deciding what in the world to attend to and what to ignore" [15]. Formal Concept Analysis formalizes the concepts concept, concept extension, concept intension, and conceptual hierarchy. We discuss this ontological commitment of FCA along two lines: a definition of concept given in a philosophical lexicon, and the international standard ISO 704.

> **Concept.** A concept is the most basic unit of thought, in contrast to judgment and conclusion, which are forms of thought composed of concepts. While a judgment makes an assertion about an issue, a concept is a notional, i. e., abstract–mental, representation of its 'whatness'; it captures an object based on 'what' it is, without already making an assertion about it. [...] For each concept one distinguishes its *intension* and *extension*. The intension of a concept comprises all attributes thought with it, the extension comprises all objects for which the concept can be predicated. In general, the richer the intension of a concept is, the lesser is its extension, and vice versa. [10, p. 39f; translated by the author]

This lexicon entry reflects a predominant understanding of concepts as being the most basic units of thought, based on which more complex entities of thought — i. e., judgments and conclusions — can be built. This understanding has grown during centuries from Greek philosophy to late Scholastic and has been stated in modern terms in the 17th century in the Logic of Port Royal [2]. It is nowadays established in the standard ISO704 [33]. The definition of formal concepts in FCA follows closely this understanding. It explicitly formalizes extension and intension of a concept, their mutual relationships, and the fact that increasing intent implies decreasing extent and vice versa. The formalization of concepts by FCA follows thus a long philosophical tradition.

The standard ISO 704 distinguishes three levels: object level, concept level, and representation level (see Figure 1). There is no immediate relationship between objects and names. This relationship is rather provided by concepts. On the concept level, the objects under discussion constitute the extension of the concept, while their shared properties constitute the intension of the concept. On the representation level, a concept is specified by a definition and is referred to by a name.[2]

While other knowledge representation formalisms like Description Logics or Conceptual Graphs mainly focus on the representation level, the focus of FCA is on the concept level. In fact, the definition of formal concepts follows closely

[2] After a discussion of the three levels, ISO 704 provides an overview over naming and definition principles, and provides quality criteria for them.

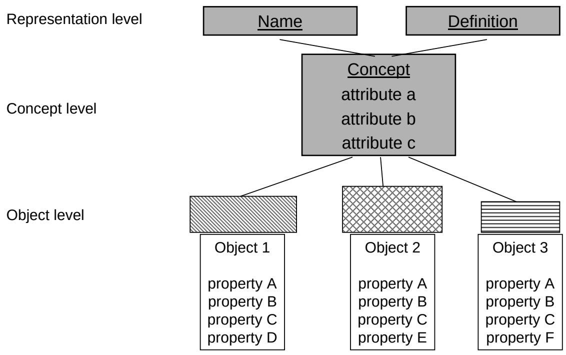

Fig. 1. Object level, concept level, and representation level according to ISO 704

the description of that level in [33]: formal concepts consist of extension and intension (only), while concept names and definitions are not within the (core) notions of FCA. Thus FCA should not be considered as competing with the other mechanisms, but rather as a complement. There is recent work following this view, for instance in combining FCA with Description Logics (e. g., [3, 60, 44, 47]) or with Conceptual Graphs (e. g., [80, 48], see also [41]) leading to the development of Contextual Logic (see Section 4.1).

2.3 Formal contexts and concepts as surrogates

"Knowledge Representation is most fundamentally a *surrogate*, a substitute for the thing itself, used to enable an entity to determine consequences by thinking rather than acting, i. e., by reasoning about the world rather than taking action in it. [. . .] Reasoning is a process that goes on internally [of a person or program], while most things it wishes to reason about exist only externally. [. . .] This unavoidable dichotomy is a fundamental rationale and role for a representation: it functions as a surrogate inside the reasoner" [15]. The authors emphasize that (human or machine) reasoning cannot deal directly with objects in the world, but only with an internal substitute: the knowledge representation.

The basic surrogates in FCA are formal contexts and concept lattices. The notion of *formal contexts* follows the understanding that one can analyze and argue only in restricted contexts, which are always subject to pre-knowledge and social conventions [80]. In applications, the transition from reality to the formal model (and back) is made explicit by the use of formal contexts; such that this interface between reality and model is always open to argumentation. Also *formal concepts*, being surrogates, only consider selected aspects of concepts, excluding for instance fuzzyness, prototypical concepts, modification over time, and so forth. In order to overcome some of the restrictions, there have been developed extensions of the formalism, for instance allowing for fuzzy concepts [43] or more expressive intensional descriptions of concepts [44, 47].

2.4 FCA as fragmentary theory of intelligent reasoning

Knowledge Representation "is a *fragmentary theory of intelligent reasoning*, expressed in terms of three components: (i) the representation's fundamental conception of intelligent reasoning; (ii) the set of inferences the representation *sanctions*; and (iii) the set of inferences it *recommends*. [...] The initial conception of a representation is typically motivated by some insight indicating how people reason intelligently, or by some belief about what it means to reason intelligently at all" [15]. The authors consider five fields which have provided notions of what constitutes intelligent reasoning: mathematical logic (e. g., Prolog), psychology (e. g., frames), biology (e. g., neural networks), statistics (e. g., bayesian networks), and economics (e. g., rational agents).

As other knowledge representation formalisms, FCA is opposed to the logistic belief that reasoning intelligently necessarily means reasoning in the fashion defined by first-order logic. The roots of FCA are best described in a philosophical view (which is close to what Davis *et al* describe as "psychological view"). It emphasizes on inter-subjective communication and argumentation, as discussed in Section 2.1. Thus — in contrast to other formalisms — FCA as such (i. e., without its extension to CKP, especially to Contextual Logic) refers the reasoning to the human user who is able to involve common sense, social conventions, views, and purposes. One of the foremost aims of FCA has always been to *support* human thinking, communication, and argumentation rather than *mechanizing* it. In [77, 81], Wille discusses the diversity in which intelligent reasoning supported by FCA takes place through sets of real-world applications. FCA in its basic form focuses on reasoning with concepts; its extension to Contextual Logic also provides a theory for reasoning about and with judgments and conclusions, including thus the triad concept–judgment–conclusion of classical philosophical logic (see Section 4.1). Reasoning with concepts comprises for instance implicational theories [20, 73, 67], clauses [24], and hypothesis generation [21].

2.5 Efficient computation within FCA

Knowledge Representation "is a *medium for pragmatically efficient computation*, i. e., the computational environment in which thinking is accomplished. One contribution to this pragmatic efficiency is supplied by the guidance a representation provides for organizing information so as to facilitate making the recommended inferences" [15]. Davis *et al* stress the importance of having a description of a useful way to organize information which allows for suggesting reasoning mechanisms and for facilitating their execution. Even though automatic reasoning is less in the heart of FCA as it is in most other knowledge representation formalisms, the question how to organize information is important for supporting human reasoning.

In FCA, information is organized in lattices. Lattices provide a clear structure for knowledge representation, which most fundamentally comprises a partial order. Unlike other partial orders (e. g., trees), they allow for multiple inheritance, which often supports a more structured representation and facilitates retrieval of the stored information. Additionally, knowledge representation in lattices is equivalent to apparently unrelated representations such as implications and closure operators. This allows to transfer knowledge into multiple formats each of which is best fit to the actual task. Last but not least, (concept) lattices are equipped with an algebraic structure (stemming from the existence of unique greatest common sub- and least common super-concepts, similar to greatest common divisors and least common multiples for natural numbers) which allows for computation within the lattice structure. As mentioned in Section 2.2, most concept lattice constructions and decompositions have as counterpart a context construction. As formal contexts are only 'logarithmic in size' compared to the concept lattice, they can be seen as a medium of efficient computation.

One can thus exploit the wealth of results of lattice theory for efficient computation. For instance, properties of closure systems are used for computing the concept lattice (e. g., [20, 68]) and valid implications (e. g., [20]); and lattice constructions are used for the efficient visualization by nested line diagrams (e. g., [76, 59]). Results from lattice theory have also been exploited for data mining tasks, for instance for conceptual clustering (e. g., [57, 40, 68]), and for association rule mining (e. g., [67]). There is still a huge open scientific potential in bringing together structural–mathematical aspects (here especially from FCA) and procedural–computational aspects from computer science.

Having discussed the attractiveness of FCA as a knowledge representation method for computer science, we will study in the next section why and how mathematics-based FCA researchers got attracted by computer science.

3 Off to New Shores

As concepts are the most basic units of thought, it is not surprising that they became important building blocks in Artificial Intelligence (AI) research. Their appearance is prevailing in Knowledge Representation (e. g., in semantic networks, conceptual graphs, description logics), but they also appear for instance in Machine Learning (e. g., in conceptual clustering, concept learning). All these approaches focus on other aspects of concepts, leading to different formalizations.

Formal Concept Analysis arose independently of the formalisms mentioned above. Integrating several ideas from quite different domains (e. g., [7, 4, 29, 16]), FCA was introduced in 1979 by R. Wille as a *mathematical* theory, in order to "restructure lattice theory", following Hentig's restructuring program (see Section 2.1). A consequence of the aim of restructuring lattice theory was that research in the early time of FCA (1980ies and early 1990ies) mainly fell into three categories: *i*) lattice theory (e. g., lattice constructions and decompositions [75]), *ii*) qualitative data analysis (e. g., a generalized measurement theory [22]), and *iii*) applications (e. g., the analysis of surveys [36]). Of course, algorithms for computing concept lattices also were an important topic (see for instance [20]).

Until the beginning of the 1990ies, the development in AI and in FCA went on almost independently. By then, the mutual perception increased. For instance, FCA researchers got in contact with the knowledge acquisition community, and AI researchers integrated FCA in their approaches (e. g., [12]). As discussed in the previous section, FCA became attractive as an AI knowledge representation, and (as we will see below), mathematicians working on FCA got interested in AI research topics. This convergence led to the aim to establish Conceptual Knowledge Processing as an extension of FCA (see next section). In 1993, the ErnstSchröderCenter for Conceptual Knowledge Processing[3] was founded in Darmstadt to support and accompany this development. Just a year later, NaviCon GmbH[4] was founded, a spin-off of Darmstadt University of Technology offering consulting based on FCA methods and tools.

The convergence of FCA with computer science research increased significantly by the series of International Conferences on Conceptual Structures (ICCS), where FCA became a topic in 1995 [37, 58]. This conference series especially stimulated the development of Contextual Logic [79] (see Section 4.1). From 1998 on, the use of FCA for Knowledge Discovery was discussed [69], and FCA was applied for improving the efficiency of data mining algorithms [5]. Today, FCA is not only considered within AI, but also in other computer science domains, as for instance in software engineering (e. g., [52]) or database theory (e. g., [51]). FCA papers are nowadays almost exclusively presented at computer science conferences and in computer science journals. The foundation of the Research Center for Conceptual Knowledge Processing (FZBW)[5] at Darmstadt University of Technology in November 2000 also witnesses the continuous interest in this research topic.

One reason for the change of orientation of FCA (and CKP) towards computer science is certainly that, in the eyes of the mathematical community, lattice theory is an almost closed research area, where almost all important problems have been solved. Further open problems, for instance the development of good lattice drawing algorithms, are not considered as genuine mathematical problems by the majority of the mathematicians.

A more important reason for the change of orientation is the fact, that computer science is — perhaps because it is still a young discipline — in general much more open-minded to discussions such as Hentig's restructuring program than mathematics is. The relationship and the interaction between user and computer is a research domain in computer science for its own sake, and, more important still, expectations and possible consequences of computer science are discussed in public.

What are future directions of Formal Concept Analysis? We conclude this section by relating Conceptual Knowledge Processing with the growing research area of *Ontology Engineering* (see for instance [39]). We believe that nowadays FCA and (parts of) AI are closer together as they sometimes seem to be. This holds especially for the consideration of the importance of the principle of knowledge representation as a medium of human expression. Partly the remaining

[3] www.mathematik.tu-darmstadt.de/ags/esz/
[4] www.navicon.de
[5] www.fzbw.tu-darmstadt.de

difference is due only to the different language they (still) speak. In fact, the importance of this principle has increasingly been discussed in the AI community in the past few years.

Interestingly, Ontology Engineering (independently) follows a trend which also served as basis for FCA. The point is that, according to J. Habermas, *ontology*, stemming from the tradition of Greek metaphysics, is constrained to a specific relationship to the world, namely the cognitive relationship to the existing world. It does not consider the subjective nor the social world. A concept corresponding to 'ontology', which includes the relationship to the subjective and social world, as well as to the existing world, was absent in philosophy. This observation was encountered in different ways. Habermas developed his Theory of Communicative Action [26] in order to provide such a concept (see Section 2.1). Habermas' theory had strong influence on the way FCA was developed. Computer scientists, on the other hand, extended the definition of the concept 'ontology' — and adapted it in a straightforward manner directly to their own purposes (which led to many controversies with philosophers). Most popular in computer science is nowadays the definition of T. Gruber, who considers ontologies as "formal, explicit specification of a shared conceptualization" [25]. A 'conceptualization' refers to an abstract model of some phenomenon in the world by identifying the relevant concept of that phenomenon. 'Explicit' means that the types of concepts used and the constraints on their use are explicitly defined. 'Formal' refers to the fact that the ontology should be machine understandable (which excludes for instance natural language). 'Shared' reflects the notion that an ontology captures consensual knowledge, that is, it is not private to some individual, but accepted by a group.

In practice, the two approaches are not far from each other. Both FCA and Ontology Engineering emphasize the importance of an inter-subjective agreement about the conceptualization, and both claim the need of a formal specification of the model. The main difference is that, in terms of ISO 704 (see Section 2.2), FCA works mainly on the concept level, while Ontology Engineering works mainly on the representation level. I. e., FCA considers extensional and intensional aspects as equal, while Ontology Engineering emphasizes on the intensional part. As already argued in Section 2.2, these views should be understood as complementary rather than competitive. We suggest thus to integrate Formal Concept Analysis and Ontology Engineering in one unified framework. Establishing this framework and working on its details are interesting topics for future research.

4 Conceptual Knowledge Discovery and Processing

In this section, we present *Conceptual Knowledge Processing (CKP)* which arose as an extension of FCA taking into account more explicitly Davis *et al*'s fourth and fifth principles; and argue why it is a reasonable choice for a framework unifying FCA and Ontology Engineering.

4.1 Conceptual Knowledge Processing

Conceptual Knowledge Processing (CKP) has as its overall aim supporting human communication and argumentation to establish inter-subjectively assured knowledge. As a computer science theory, the task of CKP is thus to provide concept–based methods and tools for acquiring, representing, and reasoning with knowledge, and for making it available for communication purposes. We analyze how FCA (with its recent extensions) fulfills this task and how it can be complemented by Ontology Engineering in the aim of supporting Conceptual Knowledge Processing. We consider the following four categories of knowledge processing: *knowledge acquisition, knowledge representation, knowledge inference*, and *knowledge communication* [38]. We will focus on technical aspects; a reflection of the philosophical foundations of CKP can be found in [78] and [81].

Knowledge Acquisition. Knowledge Acquisition techniques (in the broader sense) can roughly be categorized in two classes: those which aim at acquiring knowledge from humans (i. e., knowledge acquisition in the narrower sense), and those which acquire knowledge out of some data (e. g., documents) in which the knowledge is encoded. As we will argue below, we do not see the two classes far from each other. The latter class is subject of the research domains Machine Learning and (more recently) Knowledge Discovery. This paper has a certain focus on the second class, and therefore devotes the entire next subsection to it. There we analyze the roles of Conceptual Knowledge Discovery and of Ontology Learning.

As for the techniques for knowledge acquisition from humans, the most prominent representative within FCA is B. Ganter's *Attribute Exploration* [20] (see also [23]). It addresses the problem of a context where the object set is not completely known a priori, or too large to be completely listed. In an interactive, iterative approach, the user has either to accept a suggested implication between the attributes (i. e., she excludes potential objects) or to provide a counter-example (i. e., she provides a (typical) object) until the concept lattice is completely determined. Concept Exploration extends this approach to situations where both the object set and the attribute set of the context are not completely known a priori or too large [35, 62]. An overview over interactive knowledge acquisition techniques based on FCA can be found in [61]. Also more informal knowledge acquisition settings within FCA aim at the specification of the formal context. In a typical data analysis scenario, the first step is to establish a formal context in cooperation with the user(s). Based on the insights gained by the resulting concept lattice, the context can be refined and modified in subsequent feedback loops.

Ontology Engineering in its turn even has its roots in the Knowledge Acquisition community. From there, it brings along methodologies for knowledge acquisition, as for instance Common–KADS [49], which is currently instantiated for ontologies in the OTK ontology development framework [55]. Recent knowledge acquisition approaches within Ontology Engineering can be classified in two groups: ontology learning and instance learning (information extraction).

The first deals with learning the ontology itself (i. e., the intensional aspect) [39], and the second with learning the assignment of instances to the concepts and relations (i. e., the extensional aspect) [27].

Like FCA, Ontology Engineering emphasizes on the importance of agreeing among the domain experts on a shared understanding of the domain. One difference is that most of the Ontology Engineering approaches base the interactive knowledge acquisition process on heuristics which allow for more flexibility than FCA approaches. In general one can conclude that Ontology Engineering provides more comprehensive support for the more informal aspects of knowledge acquisition and complements thus well with the more structure-oriented techniques of FCA which come along with stronger semantics.

Knowledge Representation. Knowledge representation with FCA has already been the overall theme of Section 2. Here we focus on its relationship to Ontology Engineering.

The choice of the formalism for representing an ontology directly influences the methods and tools to be applied; there is no language–neutral Ontology Engineering. Ontologies are described in different formalisms (e. g., description logics, conceptual graphs, frame logic), depending on the task to be solved (and on the history of the researcher working on it). As argued in Section 2.2, these formalisms complement well with FCA, and first steps have been made to set up links between the underlying theories. These links have to be strengthened and are to be exploited for establishing a comprehensive Conceptual Knowledge Processing environment. From the FCA perspective, this means to extend the scope from strongly structured to semi–structured and even unstructured data, allowing to tackle more complex tasks as, for instance, in the Semantic Web.

Knowledge Inference. The second important thread in CKP is today, beside Conceptual Knowledge Discovery, the development of *Contextual Logic* [79, 82]. Contextual Logic aims at restructuring mathematical logic, following Hentig's restructuring program, in order to overcome deficiencies of predicate logic for knowledge representation [46]. It is based on the elementary doctrines of concepts, judgments, and conclusions as discussed in classical philosophical logic. In this framework, FCA is considered as a theory for concepts, while Conceptual Graphs are building blocks for a theory for judgments and conclusions. Due to space restrictions, Contextual Logic will not be presented in detail in this paper. The interested reader is referred to [79, 80, 45, 46, 82].

Davis *et al* suggest to analyze two sets of inferences for a given knowledge representation: the set of inferences the representation sanctions, and the set of inferences it recommends. As known from other mathematics–based logics, Contextual Logic currently provides a sound and complete set of inferences, i. e., a set of inferences the representation sanctions. The choice of the inferences to be applied is left to the user; Contextual Logic aims to support the user in this task by providing graphical user interfaces [18].

Ontology Engineering tools in general make use of sanctioned inferences, too, for instance for checking the consistency of the ontology, and for deriving knowledge which is not explicitly encoded. As there is no language–neutral representation of an ontology, each Ontology Engineering tool has to provide an implementation of an inference mechanism applicable to the language it uses. Additionally to the set of sanctioned inferences, Ontology Engineering tools often make extensive use of heuristics, which can be seen as implementations of sets of recommended inferences. A tighter interweaving of heuristics–based approaches with FCA and Contextual Logic is an interesting topic for future research.

Knowledge Communication. For Formal Concept Analysis, the importance of knowledge communication has already been discussed in Section 2.1. This aspect has been the driving force for the development of several tools, e. g., ConImp [11], GALOIS [12], the management system TOSCANA for Conceptual Information Systems [72] with various extentions (e. g., [70, 65, 18, 30, 71]) and the analysis tool CERNATO[6].

Ontologies also have as primary focus the support of human (and human–computer) communication. They are applied for instance for community building [53], for knowledge management [1, 55], and in the Semantic Web [6]. The Semantic Web aims at providing automated Web services based on formal knowledge representations. In this scenario, ontologies are used for instance in semantics–based portals [56, 54, 34] and for the communication of (software) agents [28].

Systems like the RFCA system for browsing rental advertisements on the WWW [13] or the Conceptual Email Manager [14] are first prototypes integrating both FCA and ontologies. The next step will be to establish interfaces between the two research and software projects 'Tockit — Framework for Conceptual Knowledge Processing'[7] and 'KAON — Karlsruhe Ontology and Semantic Web Tool Suite'[8] in order to obtain a large, stable platform for future developments.

4.2 Conceptual Knowledge Discovery

The aim of *Knowledge Discovery in Databases* (*KDD*) is to support human analysts in the overall process of discovering valid, implicit, potentially useful and ultimately understandable information in databases. The volume "Advances in Knowledge Discovery and Data Mining" [19] emphasizes that this iterative and interactive process between a human and a database may strongly involve background knowledge of the analyzing domain expert.[9] In particular, R. S. Brachman and T. Anand [8] argue in favor of a more human-centered approach to

[6] http://www.navicon.de/deutsch/sit_f.htm
[7] http://tockit.sourceforge.net/
[8] http://kaon.semanticweb.org/
[9] Following [19], we understand KDD as the overall discovering process; while *data mining* is considered as one step of KDD, namely the application of algorithms for extracting patterns from the data.

knowledge discovery ("data archeology", [9]) referring to the constitutive character of human interpretation for the discovery of knowledge and stressing the complex, interactive process of KDD as being led by human thought. Following Brachman and Anand, *Conceptual Knowledge Discovery (CKDD)* pursues a human-centered approach to KDD based on a comprehensive notion of knowledge as a part of human thought and argumentation [69, 31]. This view leads to a modified definition of what knowledge discovery is: we understand (conceptual) knowledge discovery as "information discovery combined with knowledge creation where the combination is given by turning discovered information into created knowledge" [83]. A more detailed discussion of this understanding along a list of requirements for knowledge discovery environments provided in [8] can be found in [69]. CKDD applications are presented in [63, 64, 31, 67, 68, 17].

The human–centered approach of CKDD indicates the need to distribute the work between data mining algorithms on the one hand and the user on the other hand. Ontology Learning, the knowledge discovery part of Ontology Engineering, also follows this paradigm: A. Mädche considers the process of Ontology Learning as a semi-automatic process with human intervention, since completely automatic knowledge acquisition is an unrealistic vision (today) [39, p. 52]. The approach allows the integration of a multitude of disciplines (e. g., machine learning, natural language processing, human–computer interaction) in order to facilitate the semi–automatic construction of ontologies. Instance learning, as discussed in the previous subsection, is today more based on user-centered, interactive techniques (that is why we discussed it under the heading 'knowledge acquisition' above, and not here). However, we expect that instance learning will make a more extensive use of data mining techniques in the near future.

As discussed above, we want to integrate Ontology Engineering into Conceptual Knowledge Processing. For Conceptual Knowledge Discovery, this means that Ontology Learning, Instance Learning, and FCA–based knowledge discovery should be brought together. Our vision for future research is to interweave these approaches, and to apply them for concept-based knowledge discovery. This is especially promising in the upcoming Semantic Web, where first steps towards *Semantic Web Mining* have been done [66].

5 Outlook

In this paper, we have discussed the turn of FCA towards computer science. We have analyzed why FCA is considered as a knowledge representation method within computer science, and how and why mathematics–based FCA researchers became attracted by computer science. We presented Conceptual Knowledge Processing and Conceptual Knowledge Discovery as steps in that development, and argued for a future integration with Ontology Engineering. We strongly believe that there remains a huge scientific potential in the exploitation of bringing together mathematical–structural results (especially from FCA) and procedural aspects, which will further enhance the state of the art in computer science.

Acknowledgements

I am grateful to Susanne Prediger for intensive discussions about the vision described in this paper, and to my colleagues for pointing out specific relationships to Ontology Engineering.

References

1. A. Abecker, A. Bernandi, K. Hinkelmann, O. Kühn, M. Sintek: Towards a technology for organizational memories. *IEEE Intelligent Systems and Their Applications* **13**(3), 1998, 40–48
2. A. Arnaud, P. Nicole: *La logique ou l'Art de penser.* Abraham Wolfgang, Amsterdam 1685.
3. F. Baader: Computing a minimal representation of the subsumption lattice of all conjunctions of concept defined in a terminology. In: G. Ellis, R. A. Levinson, A. Fall, V. Dahl (eds.): *Proc. Intl. KRUSE Symposium*, August 11–13, 1995, UCSC, Santa Cruz 1995, 168–178
4. M. Barbut, B. Monjardet: *Ordre et classification, Algèbre et Combinatoire.* 2 tomes. Paris, Hachette 1970
5. Y. Bastide, R. Taouil, N. Pasquier, G. Stumme, L. Lakhal: Mining Frequent Patterns with Counting Inference. *SIGKDD Explorations* **2**(2), Special Issue on Scalable Algorithms, 2000, 71–80
6. T. Berners-Lee, J. Hendler, O. Lassila: The Semantic Web. *Scientific American* **284**(5), May 2001, 34–43
7. G. Birkhoff: *Lattice Theory.* 1st edition. Amer. Math. Soc. Coll. Publ. **25**, Providence, R. I., 1940
8. R. J. Brachman, T. Anand: The process of Knowledge Discovery in Databases. In [19], 37–57
9. R. J. Brachman, P. G. Selfridge, L. G. Terveen, B. Altman, A. Borgida, F. Helper, T. Krk, A. Lazar, D. L. McGuiness, L. A. Resnick: Integrated support for data archeology. *Intl. J. of Intelligent and Cooperative Information Systems* **2** (1993), 159–185
10. W. Brugger: *Philosophisches Wörterbuch* Herder, Freiburg 1976
11. P. Burmeister: Programm zur Formalen Begriffsanalyse einwertiger Kontexte. TH Darmstadt 1987
12. C. Carpineto, G. Romano: GALOIS: An Order-Theoretic Approach to Conceptual Clustering. *Machine Learning.* Proc. ICML 1993, Morgan Kaufmann Prublishers 1993, 33–40
13. R. Cole, P. Eklund: Browsing Semi-Structured Web Texts Using Formal Concept Analysis. In: H. Delugach, G. Stumme (eds.): *Conceptual Structures: Broadening the Base.* Proc. ICCS '01. LNAI **2120**, Springer, Heidelberg 2001, 319–332
14. R. Cole, G. Stumme: CEM – a conceptual email manager. In: B. Ganter, G. W. Mineau (eds.): *Conceptual Structures: Logical, Linguistic, and Computational Issues.* Proc. ICCS '00. LNAI **1867**. Springer, Heidelberg 2000, 438–452
15. R. Davis, H. Shrobe, P. Szolovits: What is a knowledge representation? *AI Magazine* **14:1** (1993), 17–33.
16. Deutsches Institut für Normung: *Begriffe und Benennungen – Allgemeine Grundsätze.* DIN 2330. 1993

17. V. Duquenne, C. Chabert, A. Cherfouh, J.–M. Delabar, A.–L. Doyen, D. Pickering: Structuration of phenotypes/genotypes through Galois lattices and implications. In: E. M. Nguifo, V. Duquenne, M. Liquiere (eds.): *Proc. ICCS-2001 Intl. Workshop on Concept Lattices–Based Theory, Methods, and Tools for Knowledge Discovery in Databases*, Stanford, July 2001, 21–34
18. P. Eklund, B. Groh, G. Stumme, R. Wille: A Contextual-Logic Extension of TOSCANA. In: B. Ganter, G. W. Mineau (eds.): *Conceptual Structures: Logical, Linguistic, and Computational Issues.* Proc. ICCS '00. LNAI **1867**, Springer, Heidelberg 2000, 453–467
19. U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, R. Uthurusamy (eds.): *Advances in Knowledge Discovery and Data Mining.* AAAI/MIT Press, Cambridge 1996.
20. B. Ganter: Algorithmen zur Formalen Begriffsanalyse. In: B. Ganter, R. Wille, K. E. Wolff (eds.): *Beiträge zur Formalen Begriffsanalyse*, B. I.–Wissenschaftsverlag, Mannheim 1987, 241–254
21. B. Ganter, S. O. Kuznetsov: Formalizing Hypotheses with Concepts. In: Ganter, B., Mineau, G. (Eds.): *Conceptual Structures: Logical, Linguistic and Computational Issues.* LNAI 1867.Springer, Berlin-Heidelberg-New York 2000, 342–356
22. B. Ganter, J. Stahl, R. Wille: Conceptual measurement and many-valued contexts. In: W. Gaul, M. Schader (eds.): *Classification as a tool of research.* North-Holland, Amsterdam 1986, 169–176
23. B. Ganter, R. Wille: *Formal Concept Analysis: Mathematical Foundations.* Springer, Heidelberg 1999.
24. B. Ganter, R. Wille: Contextual Attribute Logic. In: W. Tepfenhart, W. Cyre (eds.): *Conceptual Structures: Standards and Practices.* LNAI **1640**. Springer, Heidelberg 1999, 377–388
25. T. Gruber: Towards principles for the design of ontologies used for knowledge sharing. *Intl. J. of Human and Computer Studies* **43**(5/6), 1994, 907–928
26. J. Habermas: *Theorie des kommunikativen Handelns.* Suhrkamp, Frankfurt 1981
27. S. Handschuh, S. Staab: Authoring and Annotation of Web Pages in CREAM. *Proc. World–Wide Web Conference (WWW 11)*, 2002
28. J. Hendler: Agents and the Semantic Web. *IEEE Intelligent Systems* **16**(2),2001, 30–37
29. H. von Hentig: *Magier oder Magister? Über die Einheit der Wissenschaft im Verständigungsprozess.* 1. Aufl., Suhrkamp, Frankfurt 1974
30. J. Hereth, G. Stumme: Reverse Pivoting in Conceptual Information Systems. In: H. Delugach, G. Stumme (Eds.): *Conceptual Structures: Broadening the Base.* Proc. ICCS '01. LNAI **2120**, Springer, Heidelberg 2001, 202–215
31. J. Hereth, G. Stumme, R. Wille, U. Wille: Conceptual Knowledge Discovery and Data Analysis. In: B. Ganter, G. W. Mineau (eds.): *Conceptual Structures: Logical, Linguistic, and Computational Issues.* Proc. ICCS '00. LNAI **1867**, Springer, Heidelberg 2000, 421–437
32. D. Horster: Habermas, Jürgen. In: B. Lutz (ed.): *Metzler Philosophen Lexikon. Von den Vorsokratikern bis zu den Neuen Philosophen.* Metzler, Stuttgart–Weimar 1995, 335–341
33. International Organization of Standardization: *ISO 704. Terminology Work — Principles and Methods.* 2000
34. M. Jarke, R. Klemke, A. Nick: Broker's lounge — an environment for multi-dimensional user-adaptive knowledge management. *Proc. 34th Hawaii Intl. Conf. on System Sciences (HICSS-34)*, 2001, 83
35. U. Klotz, A. Mann: *Begriffexploration.* Diplomarbeit, TH Darmstadt 1988

36. W. Kollewe: Evaluation of a survey with methods of formal concept analysis. In: O.Opitz (ed.): *Conceptual and numerical analysis of data.* Springer-Verlag, Berlin-Heidelberg 1989, 123–134
37. F. Lehmann, R. Wille: A triadic approach to formal concept analysis. In: G. Ellis, R. Levinson, W. Rich, J. F. Sowa (eds.): *Conceptual structures: applications, implementation and theory.* LNAI **954**. Springer, Berlin–Heidelberg–New York 1995, 32–43
38. P. Luksch, R. Wille: A mathematical model for conceptual knowledge systems. In: H.-H. Bock, P. Ihm (eds.): *Classification, data analysis, and knowledge organization.* Springer, Heidelberg 1991, 156–162
39. A. Mädche: *Ontology Learning for the Semantic Web.* PhD thesis, Universität Karlsruhe. Kluwer, Dordrecht 2002
40. G. Mineau, R. Godin: Automatic Structuring of Knowledge Bases by Conceptual Clustering. *IEEE Transactions on Knowledge and Data Engineering* **7**(5),1995, 824–829
41. G. Mineau, G. Stumme, R. Wille: Conceptual Structures Represented by Conceptual Graphs and Formal Concept Analysis. In: W. Tepfenhart, W. Cyre (eds.): *Conceptual Structures: Standards and Practices.* Proc. ICCS '99. LNAI **1640**. Springer, Heidelberg 1999, 423–441
42. Ch. S. Peirce: *Collected Papers.* Harvard University Press, Cambridge 1931–35
43. S. Pollandt: *Fuzzy Begriffe: Formale Begriffsanalyse von unscharfen Daten.* Springer, Berlin-Heidelberg 1997
44. S. Prediger: Logical scaling in Formal Concept Analysis. In: D. Lukose, H. Delugach, M. Keeler, L. Searle, J. F. Sowa (eds.): *Conceptual structures: fulfilling Peirce's dream.* LNAI **1257**. Springer, Heidelberg 1997, 332–341.
45. S. Prediger: *Kontextuelle Urteilslogik mit Begriffsgraphen. Ein Beitrag zur Restrukturierung der mathematischen Logik.* PhD thesis, TU Darmstadt. Shaker Verlag, Aachen 1998
46. S. Prediger: Mathematische Logik in der Wissensverarbeitung. Historisch–philosophische Gründe für eine Kontextuelle Logik. *Math. Semesterberichte* **47**(2), 2000, 165–191
47. S. Prediger, G. Stumme: Theory-Driven Logical Scaling. In: E. Franconi et al (eds.): *Proc. 6th Intl. Workshop Knowledge Representation Meets Databases.* CEUR Workshop Proc. **21**, 1999. Also in: P. Lambrix et al (eds.): *Proc. Intl. Workshop on Description Logics (DL '99).* CEUR Workshop Proc. **22**, 1999 (http://CEUR-WS.org/Vol-21)
48. S. Prediger, R. Wille: The lattice of concept graphs of a relationally scaled context. In: W. Tepfenhart, W. Cyre (eds.): *Conceptual Structures: Standards and Practices.* LNAI **1640**. Springer, Heidelberg 1999, 401–414
49. G. Schreiber, H. Akkermans, A. Anjewierden, R. de Hoog, N. R. Shadbolt, W. Van de Velde, B. Wielinga: *Knowledge Engineering and Management.* MIT Press 2000
50. E. Schröder: *Algebra der Logik I, II, III.* 1890, 1891, 1895. Thoemmes Press, Bristol 2001
51. I. Schmitt, G. Saake: Merging inheritance hierarchies for database integration. *Proc. 3rd IFCIS Intl. Conf. on Cooperative Information Systems*, New York City, Nework, USA, August 20-22, 1998, 122–131
52. G. Snelting: Reengineering of Configurations Based on Mathematical Concept Analysis. *ACM Transactions on Software Engineering and Methodology* **5**(2), 1996, 146–189

53. S. Staab, J. Angele, S. Decker, M. Erdmann, A. Hotho, A. Mädche, R. Studer, Y. Sure: Semantic Community Web Portals. *Proc. 9th World Wide Web Conference (WWW 9).* Amsterdam 2000, 473–491
54. S. Staab, A. Mädche: Knowledge Portals — Ontologies at Work. *AI Magazine* **21**(2), 2001
55. S. Staab, H.–P. Schnurr, R. Studer, Y. Sure: Knowledge Processes and Ontologies. *IEEE Intelligent Systems* **16**(1), 2001
56. N. Stojanovic, A. Mädche, S. Staab, R. Studer, Y. Sure: SEAL — A Framework for Developing SEmantic portALs. In: *Proc. 1st Intl. Conf. on Knowledge Capture (K–CAP '01).* ACM Press, New York 2001, 155–162
57. S. Strahringer, R. Wille: Conceptual clustering via convex-ordinal structures. In: O. Opitz, B. Lausen, R. Klar (eds.): *Information and Classification.* Springer, Berlin-Heidelberg 1993, 85–98
58. G. Stumme: Knowledge Acquisition by Distributive Concept Exploration. In: G. Ellis, R. A. Levinson, W. Rich, J. F. Sowa (eds.): *Suppl. Proc. of the Third International Conference on Conceptual Structures*, Santa Cruz, CA, USA, August 1995, 98–111
59. G. Stumme: Local Scaling in Conceptual Data Systems. In: P. W. Eklund, G. Ellis, G. Mann (eds.): *Conceptual Structures: Knowledge Representation as Interlingua.* Proc. ICCS '96. LNAI **1115**, Springer, Heidelberg 1996, 308–320
60. G. Stumme: The Concept Classification of a Terminology Extended by Conjunction and Disjunction. In: N. Foo, R. Goebel (eds.): *PRICAI '96: Topics in Artificial Intelligence.* Proc. PRICAI '96. LNAI **1114**, Springer, Heidelberg 1996, 121–131
61. G. Stumme: Exploration tools in Formal Concept Analysis. In: *Ordinal and Symbolic Data Analysis.* Studies in classification, data analysis, and knowledge organization **8**, Springer, Heidelberg 1996, 31–44
62. G. Stumme: *Concept Exploration — Knowledge Discovery in Conceptual Knowledge Systems.* PhD thesis, TU Darmstadt. Shaker, Aachen 1997
63. G. Stumme: Exploring Conceptual Similarities of Objects for Analyzing Inconsistencies in Relational Databases. *Proc. Workshop on Knowledge Discovery and Data Mining, 5th Pacific Rim Intl. Conf. on Artificial Intelligence.* Singapore, Nov. 1998, 41–50
64. G. Stumme: Dual Retrieval in Conceptual Information Systems. In: A. Buchmann (ed.): *Datenbanksysteme in Büro, Technik und Wissenschaft.* Proc. BTW '99. Springer, Heidelberg 1999, 328–342
65. G. Stumme: Conceptual On-Line Analytical Processing. In: K. Tanaka, S. Ghandeharizadeh, Y. Kambayashi (eds.): *Information Organization and Databases.* Chpt. 14. Kluwer, Boston–Dordrecht–London 2000, 191–203
66. G. Stumme, A. Hotho, B. Berendt (eds.): *Semantic Web Mining.* Proc. of the Semantic Web Mining Workshop of the 12th Europ. Conf. on Machine Learning (ECML'01) / 5th Europ. Conf. on Principles and Practice of Knowledge Discovery in Databases (PKDD'01), Freiburg, September 3rd, 2001
67. G. Stumme, R. Taouil, Y. Bastide, N. Pasquier, L. Lakhal: Intelligent Structuring and Reducing of Association Rules with Formal Concept Analysis. In: F. Baader. G. Brewker, T. Eiter (eds.): *KI 2001: Advances in Artificial Intelligence.* Proc. KI 2001. LNAI **2174**, Springer, Heidelberg 2001, 335–350
68. G. Stumme, R. Taouil, Y. Bastide, N. Pasqier, L. Lakhal: Computing Iceberg Concept Lattices with Titanic. *J. on Knowledge and Data Engineering*, 2002 (in press)

69. G. Stumme, R. Wille, U. Wille: Conceptual Knowledge Discovery in Databases Using Formal Concept Analysis Methods. In: J. M. Żytkow, M. Quafofou (eds.): *Principles of Data Mining and Knowledge Discovery.* Proc. PKDD '98, LNAI **1510**, Springer, Heidelberg 1998, 450–458
70. G. Stumme, K. E. Wolff: Computing in Conceptual Data systems with relational structures. *Proc. Intl. Conf. on Knowledge Retrieval, Use, and Storage for Efficiency*, Vancouver, Canada, 11.–13. 8. 1997, 206–219
71. *The ToscanaJ-Project: An Open-Source Reimplementation of TOSCANA.* http://toscanaj.sourceforge.net
72. F. Vogt, R. Wille: TOSCANA – A graphical tool for analyzing and exploring data. In: R. Tamassia, I. G. Tollis (eds.): *GraphDrawing '94.* LNCS **894**. Springer, Heidelberg 1995, 226–233
73. M. Wild: Computations with finite closure systems and implications. In: D.–Z. Du, M. Li (eds.): *Computing and combinatorics.* LNCS 959. Springer, Berlin-Heidelberg 1995, 111–120
74. R. Wille: Restructuring lattice theory: an approach based on hierarchies of concepts. In: I. Rival (ed.): *Ordered sets.* Reidel, Dordrecht–Boston, 445–470
75. R. Wille: Subdirect decomposition of concept lattices. *Algebra Universalis* **17**, 1983, 275-287
76. R. Wille: Line diagrams of hierarchical concept systems. *Int. Classif.* **11**, 1984, 77–86
77. R. Wille: Bedeutungen von Begriffsverbänden. In: B. Ganter, R. Wille, K. E. Wolff (eds.): *Beiträge zur Begriffsanalyse.* B.I.–Wissenschaftsverlag, Mannheim 1987, 161–211
78. R. Wille: Plädoyer für eine philosophische Grundlegung der Begrifflichen Wissensverarbeitung. In: R. Wille, M. Zickwolff (eds.): *Begriffliche Wissensverarbeitung — Grundfragen und Aufgaben.* B.I.–Wissenschaftsverlag, Mannheim 1994, 11–25
79. R. Wille: Restructuring mathematical logic: an approach based on Peirce's pragmatism. In: A. Ursini und P. Agliano (eds.): *Logic and Algebra.* Marcel Dekker, New York 1996, 267-281
80. R. Wille: 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**. Springer, Heidelberg 1997, 290–303
81. R. Wille: Conceptual landscapes of knowledge: a pragmatic paradigm for knowledge processing. In: W. Gaul, H. Locarek-Junge (eds.): *Classification in the Information Age.* Springer, Heidelberg 1999, 344–356
82. R. Wille: Contextual logic summary. In: G. Stumme (ed.): *Working with Conceptual Structures.* Suppl. Proc. ICCS 2000. Shaker, Aachen 2000, 265–276
83. R. Wille: Why can concept lattices support knowledge discovery in databases? In: E. M. Nguifo, V. Duquenne, M. Liquiere (eds.): *Concept Lattice-based theory, methods and tools for Knowledge Discovery in Databases.* Proc. of Workshop of the 9th Intl. Conf. on Conceptual Structures (ICCS '01). http://CEUR-WS.org/Vol-42/
84. R. Wille, M. Zickwolff (eds.): *Begriffliche Wissensverarbeitung — Grundfragen und Aufgaben.* B. I.–Wissenschaftsverlag, Mannheim 1994
85. M. Zickwolff: *Begriffliche Wissenssysteme in der Künstlichen Intelligenz.* FB4-Preprint 1506, TH Darmstadt 1992