

Knowledge in Information Technology

Wolfgang Bibel*

Darmstadt University of Technology

Abstract

This paper focuses on knowledge systems (also called knowledge-based – or expert – systems), their technology, history, state of the art, research problems, and great potential in future Information Technology to the benefit of humankind. Knowledge systems, a technology evolved in the discipline of Artificial Intelligence (AI), are qualitatively and conceptually quite different from the familiar data processing systems. We argue that coping with the complexity of our world and of the problems we are facing (eg. the consequences of the ongoing climate change) will fall short of success unless knowledge systems will play a much greater role than they do today. Knowledge systems are also crucial in modeling and theory formation of complex phenomena, a task which in the sciences and humanities of our days has been somewhat neglected in favor of just massively collecting data.

Nam et ipsa scientia potestas est.

Francis Bacon (1561-1626)

1 Introduction

Which intellectual achievements are impressing you most? Technological achievements like the Øresund bridge connecting Denmark and Sweden? Fiction like the Ulysses of James Joyce? Discoveries like Einstein's relativity theories? Of course, it is hard to commit with a strict hierarchy among the kinds of human achievements. But most people share a particularly high respect for *mathematical* talents and their discoveries. In fact, mathematics is commonly considered the queen of sciences reflecting highest appreciation. Wouldn't you then be impressed by a system which in a certain respect outperformed the world-best mathematicians of a whole century?

In December 1996 the system Otter/EQN by Bill McCune did just that by finally discovering a proof for a longstanding mathematical conjecture [McC96] which has withstood the efforts of a whole generation of the very best mathematicians including such heroes as the famous logician Alfred Tarski. The conjecture is that a Boolean algebra can be characterized by a certain single axiom found by Robbins or, in other words, that *Robbins*

*Also affiliated with the University of British Columbia. Author's e-address: Bibel@gmx.net

algebras – characterized by that axiom – *are boolean*. It was posed in the thirties by the Australian mathematician Robbins, and numerous attacks to prove it have been made by him and many others for more than sixty years, yet without success until Otter’s spectacular achievement. Otter/EQN is a knowledge system (KS) which may be filled with knowledge such as axioms and rules in algebra and otherwise is able to do logical reasoning about such knowledge in an extremely efficient way. KSs such as Otter with a particular emphasis on their reasoning performance are also called theorem provers.

In May 2003, to continue with a different story, the six months pilot phase of the project Halo was completed with remarkable success. The knowledge systems implemented in this short first phase were capable of answering novel questions in Advance Placement (AP) chemistry and of providing readable, domain-appropriate justifications for those answers. AP is the nationwide examination for US undergraduate students in the sciences. The best of the implemented systems in terms of grades received for its answers did (slightly) better in this examination than the average of all the participating human students.

The project Halo is a staged, long-term research and development initiative of Vulcan Inc., parent company for Microsoft cofounder Paul G. Allen’s investments and philanthropy, that aims to develop a “Digital Aristotle”, an application capable of answering novel questions and solving advanced problems in a wide range of scientific disciplines (see www.projecthalo.com for details). In 2004 Halo’s second phase extending for 30 months was announced in which the three best-performing teams from the pilot phase again compete with each other for maximal success.

Knowledge systems of a more mundane kind than the two just mentioned are in daily use by the tens of thousands worldwide. For instance, Siemens has developed a KS named GeneSim which on the basis of knowledge about patients and with access to the information in dozens of medical databases connected with the system through the worldwide web (WWW) generates diagnoses and therapies for those patients along with details on the relevant scientific literature in each case [Sch07]. Note that GeneSim achieves far more than a search engine such as Google’s which would just deliver thousands of references to text bodies, but nothing like a coherent diagnosis or therapy. Currently GeneSim aims at cancer, neurodegenerative diseases like Alzheimer and cardio-circulatory diseases.

These three examples of KSs are selected just to give an idea about the kind, and the breadth of range, of possible applications of knowledge systems technology which is the main topic of this paper. We briefly introduce into the characteristics of this technology and illustrate its great potential not only for particular applications, but for solving some of the most urgent problems of humankind such as coping with the imminent climate change.

The paper starts in the subsequent section with some preliminaries concerning the notion of knowledge, its representation and the operations involved in its processing. It proceeds in Section 3 with an account of knowledge systems and their technology. Section 4 provides the reasons why this technology has a great potential to cope with the complexity of our world and of the problems we are currently facing. It also outlines its important role in the formation of models and theories out of the data which are now

available in massive quantities. Namely, such massive data are of rather limited use if not abstracted to some kind of theory which may be grasped by humans to support our understanding of complex phenomena, a fact which tends to be neglected in the sciences and humanities of our days. Our illustrating examples range from physics over climate research to the humanities and social sciences. The paper ends with some conclusions drawn from these perspectives for the benefit of humankind.

2 Knowledge, its representation and operations

The concept of *knowledge* is not easy to define in a precise way. It refers to the psychological or cognitive capacity of humans to store pieces of information, relate them to each other, generalize them, communicate them with other humans, and activate relevant stored information at any time to plan future actions. It is particularly this capacity which distinguishes humans from animals. Only humans can build models about their environment from such information, use it in a planful way and communicate it with others. For instance, I could describe my garden in such a detail that you could imagine it, and probably recognize it in passing by or on pictures of it. Such a model is part of my knowledge, more precisely of my *declarative* knowledge. Although such knowledge in essence is private to the knowing individual, its communicability makes it potentially common. This potential leads us to mostly forget about the psychological nature of knowledge and treat knowledge as a common good.

Declarative knowledge comprises far more than just models or descriptions of objects or scenes as in this example of my garden. Namely, from such knowledge we can infer inductively generic knowledge, for instance rules such as “birds fly” or laws such as the physical laws of falling bodies. Also abstract or structural knowledge like the theorem of Pythagoras in geometry or some protocol algorithm from internet technology belong to this category of knowledge. A characteristic feature of any such knowledge is its description in terms of a complex structure of concepts which in the area of KSs is called an *ontology*. Garden, human, animal, bird, triangle are such concepts. Some of the structural relationships among these concepts are that birds are animals, but birds are not humans, and gardens have geometric forms such as square triangle.

In addition to declarative knowledge (in psychology also referred to as *explicit memory*) there is operative knowledge (or *implicit memory*) by which our brain stores skills. For instance, when I drive a bicycle or play my violin I activate such operative knowledge. The technological analogy is the implemented program which controls a robot’s behavior. But in contrast to such a program, which is accessible to us as declarative knowledge in the form of programming code, human operative knowledge is declaratively for the most part unaccessible and thus difficult to communicate. For this reason we focus in this paper mostly on declarative knowledge, simply referred to just as knowledge.

Knowledge is an ingenious product of evolution. With it the brain is able to compress the unbelievably huge and unstructured flood of information, which continuously hits our senses, into relatively small and structured portions. Hence knowledge is a decisive tool to cope with the world’s complexity. Over more than a hundred thousand years those

portions and their structures have evolved in such a way as to improve the chances for survival of our species.

In the neurosciences first rudimentary insights into the way of processing of knowledge in the brain have been achieved. For instance, we know that the synapses play an important role in the storage of knowledge. But many details are still mostly unknown. We can therefore not expect a detailed computer modelling of those processes on the biochemical level in the near future. First steps in such a direction have been done though. For instance, researchers of the university of Nevada and of the IBM Almaden Research Labs have succeeded in simulating the electric behavior of eight millions (ie. about half) of the neurons (ie. nerve cells in the brain) of a mouse [Kol07]. The Blue Brain project, pursued by the Ecole Polytechnique Federale de Lausanne (EPFL) again in collaboration with IBM, aims at a real-time simulation (on IBM's Blue-Genes computer) of the structure and electrical behavior of the human brain. But there is a long way to go until the behavior of the hundred billions of neurons in the human brain including the important biochemical phenomena will be understood with such methods. For the time being such a modelling must thus be made on a much higher level of abstraction for which enough empirical data about observable phenomena are available.

The appropriate level naturally is the natural language level because by evolution language has become the main medium for representing and communicating knowledge. Thereby language is understood here in a very general sense which also includes formal languages, gestation, mimicry, body language, artistic expressions and so forth, although we will restrict our discussion mostly to natural and formal languages. Among the latter our main focus will be logical languages which are abstracted from natural language.

The common understanding of the term of knowledge differs from our analysis above. Generally, knowledge is understood as true beliefs shared by humanity. But this general view ignores the psychological aspects of knowledge which are discussed in our analysis. However, after abstracting from these deeper mechanisms on the language level we are back at the usual understanding of the term.

Linguistic representations of knowledge are syntactic (including acoustic) expressions which, according to what we said above, carry semantics. In [Fre92] Frege introduced the distinction between the sense (Sinn) and the reference (Bedeutung), now termed intensional and extensional semantics. For instance, the intensional semantics of "my garden" says that we are talking about gardens, one of the concepts in a complex ontology, especially focussing on that garden which is in the possession of the author while the extensional semantics refers to some particular area in a place called Günterfürst. Both semantical aspects of linguistic expressions are learned in childhood.

In effect, extensional semantics is more complicated than suggested by this example because it completely ignores the psychological or cognitive aspects of knowledge. To illustrate this problem imagine a musician playing the melody consisting of the sequence of the tones c' , c'' , and g' in some rhythm on some instrument. The extensional semantics of what I just described consists of the respective acoustic signals in a physical sense which is quite different from what the brain makes out of them. For instance, the harmonic octave is processed in the brain rather differently from an arbitrary physical interval [Lan07,

[Lan06], because the wiring of the auditory tract in the brain does in fact prefer integer relations in a way which already Pythagoras has hypothesized. In other words, if we want to be precise, we need to distinguish between the cognitive extensional semantics in our psyche and the physical extensional semantics out there in the world, not only in the case of melodies but for any linguistic expression since the cognitive experience is always a somehow processed model of reality. We refer to [Bib07] for an extended discussion of this aspect of extensional semantics and here understand the extensional semantics in its usual sense. With this simplification we can think of knowledge as represented in some language.

The power of knowledge derives from the operations which can be carried out with it. The conceptually simplest such operations are the storage of new knowledge and the retrieval of stored knowledge. Our focus in the remainder of this section will be the logical operations. For a more extended discussion of these operational aspects including the logical ones see again [Bib07].

The evolutionary role of the logical operations amounts again to a compression of information. For instance, if one recognizes an object like a bird, our logical reasoning capabilities immediately open the gate to a whole bunch of related knowledge facts which otherwise would have been unknown for the particular object under consideration. For instance, we may logically infer that the bird has feathers, can fly, builds nests, and so forth. This ability is particularly helpful in the human communication process since due to logical reasoning it can be realized with a substantially reduced amount of communicated information.

In general this cognitive phenomenon allows us to infer from the available knowledge K_0 further knowledge K_1 as in the bird example. K_0 and K_1 are thus related by some cognitive relation which we denote by \models_c , hence $K_0 \models_c K_1$. It is this relation which has evolved in an evolutionary way, but is not explicitly known to us. In a similar way language representations $K_{\ell i}$ of K_i , $i = 0, 1$, have evolved and along with them a relation \models_ℓ with which we obtain $K_{\ell 0} \models_\ell K_{\ell 1}$. We know countless examples of this relation. Logicians of more than two millenia have inductively extracted from this fund of examples a precisely defined (semantic) relation \models for an abstract form of natural language, the logical language. This is the relation we will consider in the sequel and will write simply $K_0 \models K_1$ to express the relationship among knowledge chunks expressed in some logical language. On the basis of this relation syntactic logical calculi, each defining a specific syntactic relation \vdash , have been developed which model human reasoning in a formal and mechanical way.

These brief outlines of the development of logic (see [KK84] for all the details) could lead to the false impression that this development is already successfully completed. Indeed it has been extremely successful and has led to a kernel of logic which will stay untouched also in the future. But many questions are still pending and waiting for a final answer. In particular the search for the most appropriate fine-structure of the relations \models_ℓ and \models is still going on. Also, certain aspects, while widely studied with several offered solutions, seem to have not yet led to a final solution; these include the treatment of imprecision, of probabilistic knowledge, of state changes, to mention just a few important ones.

3 Knowledge Systems

Knowledge is obviously involved in any piece of Information Technology (IT). For instance, a software program clearly features a lot of knowledge although buried somehow in the code, difficult to access for humans and inaccessible for computers. It is this fact which makes software maintenance such a costly and painful task and which complicates the lives of software users in such an unpleasant way. Standard software of this kind is *not* meant by what we denote as knowledge systems (KS).

Since the evolution of the world-wide web (WWW) in the early nineties we could observe a hype under the label of knowledge management systems (KMS) which again might be confused with knowledge systems. KMSs refer to systems which manage structured text (and related material) stored in computers for access especially from remote locations. In this case the systems store knowledge in a way which may easily be understood by humans but again is inaccessible for computers. That means the systems manage text data without the slightest “understanding” of what is written in the text. KMSs again are *not* meant by what we denote as KSs.

KSs, a technology evolved in the discipline of Artificial Intelligence (AI), rather code knowledge in a way which is accessible both to humans and computers. This is possible on the basis of the logical formalisms which were discussed in the previous section. Recall that the logical language is an abstraction from natural language. Knowledge represented in logic is therefore directly accessible to humans, possibly through rewriting it in a way familiar from natural language. But the important point is that it is also accessible to the system as if it would understand the meaning of the coded knowledge like a human being. How could this be?

Of course, KSs have no access to the semantics of their coded knowledge the way we humans have. But they model the processing of knowledge in humans on the high abstraction level of language as described in the previous section and therefore behave in a way familiar to humans. In particular, knowledge may be communicated with KSs in the way familiar to humans which may lead to a change of behavior of the system, like if I hear some news and change my plans and behavior accordingly. Note that such communication does in no way involve some knowledge-hiding programming code as in standard software. Rather the language of communication is close to natural language, and in particularly advanced KSs may even be natural language.

The architecture of a KS features first of all the *knowledge base* (KB) comprising all the facts, rules, etc. available to the system. Underlying the coded knowledge is an ontology as already mentioned in the previous section featuring all the concepts used to describe the knowledge. The coding of the knowledge is done basically in a logical language, possibly in disguise. Namely, several decades of research in the area of knowledge representation have led to refined structures of representation which support the access to, and the processing of, knowledge in the machine, an issue which was of no relevance for the logical languages as originally designed by Frege and his successors.

To get a feeling of the size of such a KB we mention that the largest KS contains a couple of millions of knowledge chunks. We are talking of CYC [Len95], a KS developed under the

direction of Doug Lenat at Austin in Texas, which represents facts and rules of common sense like: “glass is fragile”, “fragile objects break if dropped”, “coffee in a cup spills over if the cup is not held upright”, and so forth. Although CYC is truly huge already, one has estimated that people have access to dozens – if not a hundred – of millions of such knowledge chunks in their memory which is to say that even CYC is just a beginning. At this point we also remind of Halo, the project mentioned in the introduction, which is of the same nature as CYC but focussing on scientific rather than common sense knowledge. The remarkable feature of KBs is that in principle they can be merged although usually some technical problems might arise due to differing representational features involved. Once these technical problems are overcome (which seems feasible in the very near future) the vision is that we just combine – most likely over the internet – all the tens of thousands existing KBs into one single huge KS which would then approach the level of quantity of knowledge in humans.

Apart from the KB the second important module characterizing a KS is the inference engine which models the logical reasoning. There is a variety of modes of inference used in KSs: deductive, inductive, and abductive reasoning, data-mining, learning, theory formation, discovery, imprecise and probabilistic reasoning, planning, causal and temporal reasoning, problem solving, explanation, argumentation, self-reflection, and so forth. Abstractly speaking, they all realize the relationship $K_0 \models K_1$ discussed in the previous section. For any details the reader is referred to the literature (such as [RN03, BHS93]).

Any KS features two additional modules. One is called the knowledge acquisition module which takes as input new knowledge in some form, puts it into the form used in its KB and stores it there in a suitable way. Ideally, the input could be represented in any form including in natural language which is not illusionary given the excellent performance of existing natural language understanding systems. The other module is the explanation module which can, if requested, explain the details leading to the system’s conclusions. In comparison with standard software this provides real comfort to the users. For instance, if a KS such as GeneSim described in the introduction proposes a therapy for a given patient, the doctor would request an explanation why the system thinks the recommended therapy would be preferable to any possible others. Since a KS models human reasoning the doctor (or even the patient) could follow the system’s line of reasoning without any expertise in IT.

The reader might consult any textbook on KSs for further details, for instance [Ste95]. In practice, KSs typically are embedded within standard software systems so that often one is not even aware of the kind of system actually doing the job. The technique of KSs, having emerged in the 1970s, in the meantime has become rather mature, but again many improvements are still imaginable. These include the compatibility or interoperability of different KBs, the enhanced comfort in using KSs, the ease of acquiring and collecting knowledge, the efficiency of performance, the integration and possibly harmonization of different modes of reasoning, and so forth. We are even thinking of more fundamental improvements such as associating with knowledge chunks certain sensoric states experienced by the system in observing the corresponding scenario in reality. This means that a system could have a realistic “imagination” of an expression like “my garden” or, in other

words, would be empowered by some sort of a cognitive part of extensional semantics as discussed in the previous section.

For what follows in the subsequent discussion we imagine that in the near future KSs with all these wonderful features (or at least many of them) will be available. This is not as illusionary as it may sound given the advanced state of the art in KS technology as well as the rapid progress of innovation in IT. Before we discuss the great potential of KSs let us finish this section by pointing out their characteristic properties.

KSs are inherently *declarative*. This is because their central KB component contains knowledge expressed in a declarative way which may be understood immediately by any kind of human user. KSs also are *additive* which means that adding, or removing, an item of knowledge is all needed to update or modify the state of the system which otherwise need not be changed in any way. Both of these two features are in stark contrast to standard software systems which is the fundamental reason for their extreme user unfriendliness. Standard software systems are more efficient than KSs though.

4 The promise of knowledge systems

Let us start this section by pointing out that our times are characterized by a fundamental dichotomy of global “Promise or Perish”. On the one hand the general rapid progress in technology promises the solution to major problems by further dramatic increases in computational powers [Bib05b] as well as by synergies in the “converging technologies” exemplified by nano-, bio-, info-, and cogno- (NBIC) technologies [BAdC⁺04]. For instance, Ray Kurzweil writes in [Kur05]: *ITs are already deeply influential in every industry. ... in a few decades, every area of human endeavor will essentially comprise ITs and thus will directly benefit from the law of accelerating returns*. By this law he denotes the fact that in certain aspects the progress in IT according to hard historical data accelerates in an exponential way.

On the other hand there are obvious indications that humankind causes ever more serious problems which might lead to disastrous developments. The impending climate change is just one among a frightening number of such developments (extinctions of species, destruction of virgin forests, global epidemics, failure of governance, social disintegration, digital divide, war of cultures, etc.). In this section we want to outline what role KSs could play in this subtle balance.

As we pointed out in the previous section the technology of KSs has matured and is involved in tens of thousands of systems in daily use worldwide. Although this might sound impressive and satisfactory, in comparison with the millions of standard software systems in daily use the percentage of KSs is comparatively small. The bottom-line message of this paper is that this relatively small penetration of the KS technology in IT amounts to missing great opportunities in virtually any area using IT – and which area is not using it? In the present section we would like to illustrate this great potential for the benefit of our societies and of humankind.

We start with kind of an inhouse application. Anyone using a computer can tell many stories about the frustrations with the systems involved, especially once something is not going as expected. Standard software is parameter driven and in a certain sense functional. In case of need for some modification of the system this requires change of the code in order to modify the functionality or change of the parameter values. Except for a few specialists users are absolutely unable to understand the code to an extent to be able to carry out such cryptic modifications. Help functions offered as a remedy by the systems are rather more like a joke than of any help.

With the technology of KSs software could be built in a much more user-friendly way. This would involve a declarative specification of the system like in a KB. The actual code of the software could be associated with this specification through the relation defined by the synthesis of the code from the specification, ie. by the act of programming. In case of a required modification the user would declaratively specify the necessary change in terms of his or her needs which through the knowledge acquisition component would cause the resulting changes in the system specification. Through the synthesis relation this would lead to consequential changes in the actual code carried out automatically by the system. We have the technology to realize such a comfort and thus there is no excuse for not introducing such an advanced and user-friendly technology. Unfortunately, software companies seem to shy away from giving their user this enhanced freedom from their control which would obviously involve also economic consequences.

Note that in this scenario the synthesis relation is not yet required to be established automatically but could be realized as in current software practice. This is because, given the relation, it is a technically relatively easy task to compile small changes in the specification as envisioned in this scenario in the actual code. But once the synthesis relation would be made explicit in such a way, one might also expect a revival of the efforts in the area of program synthesis whereby programming code is automatically synthesized from declarative specifications in a knowledge-based way. Much of programming is relatively simple routine work which could well be carried out by computers themselves. Software companies for obvious reasons are even less interested in making this extended vision real.

A more declarative approach to programming would make a substantial contribution towards coping with the complexity in building software systems. As software is the key tool in virtually all areas of science, technology and their applications, such an enhancement in software technology would, in turn, result in a large step towards coping with the complexity of the imminent tasks at all levels and in all areas. Given this exceptionally high potential impact, it has always been a mystery to the present author why not even the IT researchers give program synthesis their highest possible priority. Admittedly, the efforts in object-oriented programming (OOP) and especially around the unified modelling language (UML) are steps into the right direction, but certainly only preliminary ones. Also logic programming amounted to a step into the right direction but for complex, mainly sociological reasons was mostly ignored by the software community.

However, there is one area in IT which is pushing vigorously towards a KS technology, namely the area of the WWW. We already mentioned at the beginning of this section the area of knowledge management systems (KMS) in this connection. It is easy to understand

why particularly in this application the lack of semantic understanding of the involved knowledge becomes painfully recognizable to the casual user as a deficiency. For instance, if you ask the WWW for the birthday date of a friend of yours or for “the man who invented the printing press”, you may get all kinds of junk information but mostly not the correct answer – or at least an apology from the system for not being able to provide the appropriate data. Anyone can see in such a strange behavior that something is badly wrong with the technology in this respect. The technological answer to this weakness is the *semantic web* which differs from the web in its present state exactly by the KS technology. We hope that through this popularisation the technology will spread out to all the other areas mentioned above and below.

Science aims at a rational understanding of the world to secure survival in the best possible way. For instance, humankind currently faces the vital question whether, and how, its activities influence the climate. Actual climate models in fact predict a change in our climate with disastrous consequences. Hence one tries to understand the climatic processes to an extent as to prevent the causes for such undesirable changes. This scientific goal is in many aspects extremely complex. Numerous components do contribute, several of them perhaps not even recognized, and the causal structure of these contributions is badly understood. In other words, we do not even have a clear specification of the task underlying that scientific goal, let alone a solution.

What we have just said of the climatic research holds for many other areas in a similar way. Further examples are the processes in the human body, or even just in a single of its trillions of cells, the brain and its phenomenological properties like intelligence and consciousness, the human psyche, the biosphere, the mechanisms at the molecular level of matter, the world economy, and many more. They are characterized by the non-linearity of the causal relationships of the system’s components, the emergence of properties through interaction of components at a lower level, and, last not least, by the lack of a clear research task specification as just illustrated for the case of the climate, leading to a “stoking in the fog” in pursuing the goal of scientifically understanding these phenomena.

The prevailing tendency in the scientific community consists in focussing mostly on the collection of data, in fact of massive amounts of data, at the expense of comparable efforts in theory formation. The technological reason for this is the availability of IT tools for automating measurements, opinion polls, etc. and for storing nearly unlimited amounts of data. Since the understanding of phenomena in essence amounts to induce theories from given data, data collecting does indeed solve part of the problem. But the theory formation part needs similar, if not greater, efforts. Especially it, like data collection, also requires IT tools. KSs offer such tools. The more data are available, the more KS technology is needed, ie. the problem is sharply aggravated by the enhanced technology for data collections. This is because grasping the structures buried in masses of data in the range of billions or even trillions obviously exceeds human intellectual capabilities. In consequence, it is high time to engage KSs for technologically supporting the challenging task of theory formation. The report [Emm06] gives an excellent impression of what we could expect in this direction. The methods currently in use in this respect described in some detail in [Bib07], such as modelling by differential equations or informal texts, are simply insufficient as argued in detail in the same reference.

Systems of the kind we have in mind would start out with a KB consisting of all knowledge available for the topic of interest including all the collected data (note that data are a primitive form of knowledge). Since from a formal point of view any KB is in effect a theory, the task consists in inducing from this initial theory by inferential methods for theory formation a more general and minimal theory covering the given data. This includes the identification and elimination of inconsistencies which might have found entry into the KB as well as the elimination of redundant data covered by the remaining theory. Of course, human guidance in this process would be substantial for success in complex topics. The resulting more compact KB could immediately be used for solving problems, explaining observed phenomena, or predicting future events or phenomena. As new data would become available through such activities they would be integrated into the KB, so that the cyclic process of theory formation would start anew. This way we would generalize the KB towards true knowledge of the kind which Bacon had in mind in the four hundred years old citation given at the outset of this paper.

This description illustrates that a characteristic feature of KBs is their usability as an anytime additive system which is crucial for coping with the “stoking in the fog” mentioned earlier. This means that at any time in the process we do have a theory applicable for problem solving, explanation and prediction and that we may use additionally available knowledge by just adding it into the KB.

Above we mentioned a number of possible subjects for the application of KBs in this manner. These include subjects from the natural sciences as well as from the social sciences and humanities. While theory formation in a mathematical precise form is standard in the natural sciences – although with methods which have reached their limits due to the amount of data to be taken into account as explained above –, the social sciences and humanities balked at the evolution of precise, formal, and experimentally testable theories. We therefore see a particularly revolutionary application of the KB technology in those “soft” disciplines like economy, society, law, etc. towards their advanced scientification. In [Bib05a] the author has outlined the potential of such technological revolution for the area of law.

Not only the sciences but also the practical applications of the daily life, be it of private persons or of entrepreneurs or politicians, will profit from the technology of KSs. In the book “Lehren vom Leben” (lessons from/about life) [Bib03] the author has described numerous examples illustrating this thesis. Not least KSs may assist individuals or organizations in predicting the consequences of their actions. By this way they may take greater responsibility for their behavior not only under the usual local considerations but also under global aspects taking the effects on others and on the environment into account.

5 Conclusions

This paper has outlined once again the attractions which are offered by the technology of knowledge systems (KS). We have pointed out that this technology has matured and is in daily use in tens of thousands of systems across the world. But in comparison with standard software systems the share of KSs is still rather small.

As the paper tried to demonstrate the reason for the relatively small penetration of this technology lies not in some weakness or hidden disadvantage. On the contrary, in comparison with standard systems KSs offer overwhelming advantages. In fact, they have the potential for coping with the complexity of the problems humankind is currently facing. The paper has described the features of KSs which nourish these hopes. It also has illustrated the range of problems which might be attacked in this manner and particularly has emphasized the great potential for the social sciences and humanities. The paper gives also a number of cues for future research in the area of KSs.

An explanation why the potential of KSs is not tapped to a much greater extent would require a sociological analysis rather than a technological one. Such an investigation would have to focus on the challenge posed by knowledge system technology to our intellectual capacities in terms of its further development as well as on traditions in our societies. As to the first, logic and inference mechanisms like any formalism in the structural sciences (ie. mathematics, informatics, intellectics, etc.) require the ability for abstract thinking which is at people's disposal only in relatively rare cases. For those it requires the furtherance and encouragement beginning in early childhood and an environment which allows for absolute concentration on risky problems over long periods of times, circumstances which are rarely provided anywhere in our world.

Concerning the societal traditions it is clear that KSs penetrate the province of capacities which distinguishes humans from animals and in which humans continue to excel. The competition in this domain through KSs meets human resistance for obvious reasons, not least economic ones as a number of professions like law or medicine fear to loose their lucrative sources.

In this situation KSs might continue to further evolve at a rather slow pace, except perhaps if, first, the semantic web discussed in the previous section will become successful and popular and, second, the problems like those caused by the climate change became so severe that they leave us no choice anymore other than grasping the opportunities of knowledge systems technology. I wished we would act by conviction at a somewhat earlier point.

Acknowledgments. I greatly appreciate the invitation by Giovanni Jacovitti, Alberto Pettorossi and the organizers to the presentation of the paper in this conference. In addition I thank Alberto Pettorossi for providing me with a number of hints for improvements of the text and Thomas Christaller and Michele Missikoff for discussions.

References

- [BAdC⁺04] Wolfgang Bibel, Daniel Andler, Olivier da Costa, Günter Küppers, and Ian D. Pearson. Converging technologies and the natural, social and cultural world. Technical report, Commission of the EU, 2004. http://europa.eu.int/comm/research/conferences/2004/ntw/pdf/sig4_en.pdf.

- [BHS93] Wolfgang Bibel, Steffen Hölldobler, and Torsten Schaub. *Wissensrepräsentation und Inferenz*. Vieweg Verlag, Braunschweig, 1993.
- [Bib03] Wolfgang Bibel. *Lehren vom Leben – Essays über Mensch und Gesellschaft*. Sozialwissenschaft. Deutscher Universitäts-Verlag, Wiesbaden, 2003.
- [Bib05a] L. Wolfgang Bibel. AI and the conquest of complexity in law. *Artificial Intelligence and Law Journal*, 12:159–180, 2005.
- [Bib05b] Wolfgang Bibel. Information technology. Technical report, European Commission, 2005. ftp://ftp.cordis.lu/pub/foresight/docs/kte_informationtech.pdf.
- [Bib07] Wolfgang Bibel. Wissenssysteme und Komplexitätsbewältigung. In Th. Leiber, editor, *Denken und Handeln in einer komplexen Welt – Festschrift zum 60. Geburtstag von Professor Klaus Mainzer*. Hirzel Verlag, Stuttgart, 2007.
- [Emm06] Stephen Emmott. Towards 2020 science. Technical report, Microsoft Research Cambridge, 2006. <http://research.microsoft.com/towards2020science/>.
- [Fre92] G. Frege. Über Sinn und Bedeutung. *Zeitschrift f. Philosophie u. philos. Kritik*, 100:25–50, 1892.
- [KK84] W. Kneale and M. Kneale. *The Development of Logic*. Clarendon Press, Oxford, 1984.
- [Kol07] Peter Koller. Superrechner simuliert halbes Mausgehirn. *Computer Zeitung*, 38(19):2, 7 May 2007.
- [Kur05] Ray Kurzweil. *The Singularity Is Near – When Humans Transcend Biology*. Viking, Penguin Group, New York NY, 2005.
- [Lan06] Gerald Langner. Oscillating neurons in the cochlear nucleus: I. Simulation results, II. Experimental basis of a simulation paradigm. *Biological Cybernetics*, 95:371–392, 2006.
- [Lan07] Gerald Langner. Die zeitliche Verarbeitung periodischer Signale im Hörsystem: Neuronale Repräsentation von Tonhöhe, Klang und Harmonizität. *Zeitschrift für Audiologie*, 46(1):8–21, 2007.
- [Len95] Douglas B. Lenat. CYC: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11):33–38, November 1995.
- [McC96] William McCune. Robbins algebras are boolean. *Association for Automated Reasoning Newsletter*, 35:1–3, 1996.
- [RN03] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall Series in Artificial Intelligence. Pearson Education, Upper Saddle River, N.J., 2nd edition, 2003.

- [Sch07] Tim Schröder. Gezielte Datensuche. *Pictures of the Future*, pages 70–72, Frühjahr 2007. ISSN 1618-548X.
- [Ste95] Mark Stefik. *Introduction to Knowledge Systems*. Morgan Kaufmann, San Francisco, 1995.