

AI and the conquest of complexity in law

L. WOLFGANG BIBEL

Darmstadt University of Technology, 64289 Darmstadt, Germany

E-mail: bibel@gmx.net

What is possible is not independent of what we believe to be possible. The possibilities of such developments in the practical world depends upon their being grasped imaginatively by the people who make the practical world work.

Neil MacCormick (1993, p.18)

Abstract. The paper identifies some of the problems with legal systems and outlines the potential of AI technology for overcoming them. For expository purposes, this outline is based on a simplified epistemology of the primary functions of law. Social and philosophical impediments from the side of the legal community to taking advantage of the potential of this technology are discussed and strategic recommendations are given.

Key words: AI and law, peddems of legal system, epistemology of law, formalized law, legal knowledge systems, legal ontologies, philosophy of legal reasoning

Several indicators point to a trend towards a more demand-driven evolution in science and technology as opposed to a technology-driven one. As research and technology become an ever more costly and complex enterprise in contrast to the individual thinkers and inventors of the past, we need to ask beforehand what the benefits and effects of some project might be. One might even start out by identifying an existing problem and only then ask what kind of research and development would be needed to solve the problem, in a way which approximates an optimal solution under social, ecological and economic considerations. This is the approach I want to take here.

I used this approach in my book “Lehren vom Leben”, i.e., lessons about/from life (Bibel 2003), which provides an analysis of all aspects of our lives and the world from the point of view of an intellectician and identifies many problems of our society which could be alleviated by AI technology. Law is one of the areas briefly discussed (in its Section 4.7.3). The present paper expands these considerations to a more thorough analysis and provides a perspective of the legal field.

Law is a timely domain to consider since, after the penetration of Information Technology (IT), including AI, into the domains of communication, media, engineering, management, administration and so forth, IT is

now knocking at the doors of the social sciences and their applications. It is time for the legal domain to make the transition from the industrial age to the age of information and knowledge. The paper outlines how such a transition might be achieved.

We begin in Section 1 with a description of a number of problems in the legal domain which came about due to societal changes. These problems demonstrate that there is indeed also a need for changes in the legal field. Section 2 then provides the author's abstract view of the law, in order to lay the basis for the perspective developed thereafter. Section 3 outlines what AI methodology and technology could provide for tasks in the legal domain, in light of the problems identified earlier. Section 4 recommends a number of concrete steps to be taken in the near future. Section 5 concludes the paper. For obvious reasons the material is rooted in the culture of the Western World (with a bias on the German and European one). I am sorry for my ignorance of the law in other cultures.

The reality of the legal domain

In order to paint a balanced picture, let us start with the ideal behind the legal system. Any society must obviously be based on a set of rules and norms which govern the behavior of its citizens. In order not to lay justice in the hands of emotionally charged fellow citizens or subject justice to the mercy of the mood of individuals, civilization has evolved the instrument of laws written down on the basis of general experience and public judgment. Individual behavior or misbehavior is then judged by applying some appropriate rule of law in a rather logical way (even though judges rarely apply logical rules in any conscious way). The idea behind this way of judging is still highly attractive and no substitute whatsoever is in sight. So far so good.

Reality, unfortunately, has made a somewhat distorted system out of this original ideal for many reasons, some of which discussed next. First of all, we have become a litigious society; the number of lawsuits each year reaches ever higher levels, notwithstanding the encouragement of arbitration and other alternative dispute resolution procedures. In Germany, the number of civil lawsuits (excluding family law disputes) filed in 2001 at the lowest court level was 1,421,404. That is 1.8% of the entire population, including babies and elderly people. Since lawsuits end with either one party being the loser or a settlement agreement in which both parties accept less than they feel they legally deserve, the increasing number of lawsuits brings with it an increasing number of frustrated citizens. Oftentimes the perception is that attorneys (derogatorily termed "pettifoggers") talk their clients into going to court, occasionally because of their economic interests.

Due to the high number of lawsuits, courts are extremely overburdened. As a consequence it takes years (in Germany typically 3 years if an appeal involved), sometimes more than a decade, until cases are finally decided, having gone through all levels. Achieving justice in the long period of a decade often amounts to the same as losing the case right away, especially in cases involving issues from evolving areas of science and technology, where scientific evidence often is not clearly understood.

Also, many judgments appear unjust to the average citizen. The legal system has departed to a certain degree from the popular ideal of justice. This is a result of our representative system of democracy, which does not guarantee that elected politicians reflect the long-term judgment of the general public.

There is a substantial lack of legal unity. Even the Supreme Court of the US, or the equivalent institution in Germany, is sometimes unable to apply general legal principles to several cases in a consistent way. The judges deciding a case more often than not come to different opinions. Different courts often decide equivalent cases in opposing ways. Oftentimes decisions depend upon the persuasiveness of an attorney's argument. While justice is blind and judges are ethically bound to be impartial, the political and moral values of judges play an important role nonetheless.

One reason for the inconsistency of judgments lies in the explosion of the number of statutes and regulations. In the US, there are more than one hundred federal regulatory agencies which together issue more than 4500 new regulations each year (Coglianese 2004). As a consequence, the number of regulations in the *Code of Federal Regulations*, the official repository of all binding federal regulations in the US, reached in 2000 a level of about 140,000. If all regulations from the states of the US are added, the figure becomes a multiple of this level. In Europe, with its even more complex legal structure, the figures are probably even higher. For instance, the current German Secretary of Finance, Hans Eichel, mentioned in a recent interview that there are already 70,000 regulations ("Verordnungen") just in the German tax code, which makes up only a small fraction of German law. In addition, European legislation also has to be taken into account, so that in total we are talking about many hundreds of thousands, perhaps even millions, of laws and regulations exerting a strong influence on our daily lives. Moreover, since judgments are made on the basis of previous court decisions in similar cases, retrieved by relatively crude information systems, judges and attorneys can be drowned in a sea of data when handling cases. It is obvious that with current technology no judge can cope with such extensive complexity. Faulty judgments are bound to be the rule rather than the exception.

The complexity is not just one of sheer numbers. Rather, regulations are related to each other in complex and intricate ways. For instance, they are ordered in a hierarchical way depending which legislative body or regulatory agency enacted them, e.g., the state legislature of a Bundesland, the

Bundestag or the Commission of the EU, to mention three prominent ones for the German situation. Also, other laws are referenced implicitly by the legal concepts appearing in any paragraph. This interrelationship has not been established in a systematic way according to a precise design plan (like, for instance, the one for the Internet) but rather, from the point of view of an IT specialist, it has grown in a rather accidental way, making any careful analysis difficult and prone to errors. Finally, there are the many unwritten norms forming the culture of our civilization and exerting a great influence on the way laws are interpreted. Enacting new regulations therefore commonly requires 3 years or more for the responsible agency (Coglianese 2004).

The consequence of this complexity is that the ordinary citizens are unable to know about their rights. Since mostly only the wealthy and well-educated can afford to hire lawyers, ordinary citizens are often subjected to unjust treatment (e.g., regarding taxation or social benefits).

A large group of problems with the legal system centers around legislation. The process of drafting laws by parliaments or regulatory agencies is everything but high quality. Legislatures react instantly to every shift in the sentiment of the electorate, rather than respecting the long-term judgment of the public. The more or less accidental compromises achieved when drafting legislation are negotiated in the course of long debates typically ending in sessions of committees lasting many hours until the negotiators are tired enough to give in. Usually the shortcomings of the results show up already the very next day, requiring many amendments in the course of time. For instance, according to a recent statement of the leader of the Free Democratic Party in Germany, Mr. Westerwelle, the law of health insurance in Germany has been amended 200 times since its original enactment. So judges have to be careful to apply the version of law which was valid at the time of the events of the case, complicating things even further.

It is clear that this occasionally erratic process of legislation does not at all fulfill the requirements expected for a task of such importance, affecting society in essential ways. These requirements include a careful analysis of the problem being discussed, the preferences of (informed) citizens, current expert opinion on the topic, general norms and political preferences, the effects of the proposed norm, the appropriate level of norm, the embedding of the norm into existing law, and many further such issues. Since none of these requirements is satisfied to a sufficient degree and monitoring compliance with existing law is hardly feasible in such a jungle of regulations, law has lost some of its normative character with dramatic consequences. One of them is the massive breaching of the law, either unintentionally due to confusion or in a determined manner. For instance, a recent study by Transparency International came to the conclusion that corruption in the health sector in Germany presumably amounts to some 20 billion Euros per year (about a 10th of the entire federal budget).

A more general problem is the fact that the jurisprudence continues to be based on what might be termed “commonsense psychology” and as a scientific enterprise tends to ignore hard scientific facts from Cognitive Science. For more details of this point see (Bibel et al. 2004).

An instance of this problem is the imprecision inherent in the terms and legal language. Legal texts are written in a semi-artificial natural language. This means that a legal text is generally more precise than one in natural language. But legal texts nevertheless suffer from many problematic features of natural language in terms of their possible interpretations or misinterpretations and are therefore subject to psychological influences.

As a final remark we mention that in comparison with the management of business and the production of goods the processes of the legal profession in the widest sense appear to be rather archaic. More specifically, professional Information and Communication Technology (ICT) support for managing legal processes in general lags behind ICT support of analogous processes in competitive corporations (just think of the degree of IT-driven automation in production lines or the degree of ICT support in modern airports, to mention just two out of many examples). But this issue is beyond the scope of our AI focus.

With all these issues we have certainly not at all exhausted the list of problems in the legal domain and the reasons for our statement that the legal system is facing problems causing enormous costs on the level of many hundreds of billions of Euros worldwide. An exhaustive treatment of such a list is not the purpose of the present paper. Nor is it my intention to blame anyone, let alone an entire professional community, for wrongdoing. Rather, the selected problems should suffice to convince even the most callous reader that the legal domain is in anything but an ideal state. My main point is that the legal system is functioning as if we were still living in the industrial age and has not aligned itself with the information age. In other words, several of these problems could be substantially improved by technological means, but neither the legal community nor the political establishment (which is strongly influenced by the former) are ready to take convincing steps into such a direction. Before we describe the technology we have in mind in this direction, we first need to clarify what kind of system the legal system is, when viewed from the perspective of *Intellectics* (i.e., Artificial Intelligence, or AI, and Cognitive Science, or CogSci).

The epistemology of law from an intellectics perspective

We are now going to give a brief sketch of the epistemology of law from our perspective. The human world modeled by legal systems may be viewed as follows. There is a set of human agents in an environment. The model we

may think of when discussing these agents is the one used in any standard AI textbook, such as (Russell and Norvig 2003). The environment consists of further biological agents such as animals, plants, etc., and the physical world around all these agents.

The human agents interact with this environment via communication (such as speech acts) and physical action, both simply termed *acts* in the following. One may regard an act as a function from states to consequences, i.e., acts change the state of the world in certain ways, whereby “world” refers both to the agents and their environment. In this paper we use the notion of an act in a very broad sense, so that both typing a single letter as well as writing an entire book, consisting of typing letters and other actions, both are considered to be acts (or processes if you prefer). On the other hand, we may consider acts to be distinct if they are carried out in different world contexts even though they are otherwise the same.

In the same vein, we generalize the notion of a (human) agent to cover also groups of agents and legal entities, such as corporations. Such groups of agents together may also be regarded as agents and may perform acts in our generalized understanding (such as manufacturing a product) which are combinations of many other acts. Therefore, unless explicitly stated otherwise, “agent” from now on will refer both to individual human agents or groups of human agents.

One basic function of legal systems is to classify a certain subset of all possible acts (or non-acts) as “illegal.” Let us refer to this as the *norm* function. Another basic function, the *sentencing* function, consists of specifying a punishment, penalty, fine or sanction (for uniformity again considered as an – enforced – act) for each “illegal” act. Last but not least there is the *fact-finding* function, which assigns an act to given evidence. Much of the primary functionality of the legal system is already covered by these three functions of norm, sentencing and fact-finding.

The purpose of the norm function is to inform the agents about how to behave. This includes, for instance, also authorizations. The function can have more than just two values, e.g., legal, undesirable, and illegal. Note that the term “illegal” is not meant to refer only to criminal law, but is used here in the most general sense covering, for instance, negligence in running schools properly, overlooking unjustified tax exemptions, bankruptcies, etc. The purpose of the sentencing function is to teach misbehaving agents a lesson (e.g., in the form of a fine) and/or preventing them from performing further illegal acts for a certain amount of time (e.g., by way of imprisonment). The purpose of the fact-finding function is to extract from the evidence provided to judges, e.g., by the plaintiff, defendant, and witnesses during trial, the likely act to which the other two functions are to be applied.

Abstracting a substantial part of the legal domain to just three functions might be regarded as an oversimplification (which it is). But with the right

abstraction essential features become visible which otherwise could remain unnoticed. I hope to demonstrate this virtue of my chosen abstraction in the following.

The next fundamental question is how these three functions can be specified. The current legal system uses for this purpose a semi-artificial natural language (and the natural logic of legal professionals). The legal code is written in this language and in some way describes the first two of the three functions so that a given case, or rather act (in our general understanding), can be judged as legal or illegal, possibly with a sentence to go along with the judgment. In the next section we will discuss a more systematic form of specification.

The legal system has of course more components than just the legal code. In particular, two main parts are the legislature and the courts. The legislature, in response to and as result of political processes, changes the legal code by enacting new laws, amending existing laws, or invalidating laws. In our abstract view, the legislature changes the specification of the norm function. The way these changes are made is subject to prescribed procedures, for voting in parliament, signing by empowered authorities, publishing in certain journals, and so forth (see the next section for further details). These procedures are also acts in our generalized sense, so that their regulation also falls into the domain of our norm function.

The courts act in response to complaints by plaintiffs or public prosecutors, who present a case and ask the court to make a judgment, i.e., to determine the values of the two basic functions. Litigation also follows prescribed procedures, and thus may be viewed as being regulated by the norm function. This includes, for instance, the requirement that judges take previous decisions in analogous precedent cases into consideration.

There are many more aspects of legal systems. For instance, there is the hierarchical structure among the various courts. The members of legislative bodies have to be elected and authorized in certain ways, judges have to be nominated, would-be legal professionals have to be educated in jurisprudence and the practice of law, attorneys must pass the bar examination, apply for admission to practice based on moral and ethical character, and then, once admitted, are bound by the ethical standards of the legal profession. Each jurisdiction may have its own ethical standards for its admitted attorneys. Also, there are mutual influences of the various parts of the system on each other, e.g., court decisions influence the legislature, judges and lawyers exert a mutual influence, the educational system influences the way judgments are reached, etc. It is not the purpose of this paper to exhaustively cover all aspects of legal systems. The book "Rechtssysteme" (Eckhoff and Kristian Sundby 1988) provides a more comprehensive view, although it is based on system theory as it was in vogue in the 70s of the last century. More modern AI contributions are discussed in a special issue of the Artificial

Intelligence Journal (Rissland et al. 2003). The forthcoming book entitled “Legal Reasoning: A Cognitive Approach to the Law” (Sartor 2005) is an excellent source for all theoretical issues.

The potential of methods and technologies from intellectics for the legal domain

So far we have presented an abstract view of legal systems and pointed out a number of problems with existing systems of this kind. The main message of the paper is the great potential of methods from Intellectics for alleviating or even overcoming some of these problems. We see this potential in the tasks of specifying the norm, sentencing and fact-finding functions, of translating political will into legal code, and of applying the norm and sentencing functions to a given case to determine the limits of judicial discretion. In order to describe this potential in more detail, we first need to understand how Intellectics would specify the norm function; this is the first topic of this section.

FORMALISING LAW

Let us first restrict our attention to the specification of the norm function, since the sentencing function can then be handled in an analogous way and the specification of the fact-finding function will be discussed shortly. Any such specification requires the domain to be represented in a formal way. As we have said, the current legal system uses a semi-artificial natural language for this purpose, together with the natural logic of legal professionals. Since we already pointed out a number of flaws in this approach, we want to do better, by using a more precise language and building the logic into the representational system, as usual in any knowledge representation system.

We are thus faced with the task of representing a potentially infinite number of acts in some formal language in a finite way. This first requires a representation of the world in a finite language, resulting of course in a huge complexity reduction. The classification would be achieved if all illegal acts are characterized in this language. The task of finitely characterizing a potentially infinite set of acts is a standard one in Computer Science (CS). Every programming language allows this to be done for subsets of the set of acts potentially executable on a computer. It is this analogy which makes CS and AI so relevant for the legal domain, since semantic differences in the nature of the acts (or processes) are irrelevant in a formalization.

There are two qualitatively different ways of specifying sets of acts. One consists in a program which specifies the acts deterministically and in considerable detail. The other only specifies the essential aspects such as

input/output behavior and possible further constraints, mostly in a logical language, while leaving details of the acts open for an appropriate choice in a particular realization. We will refer to the latter way as a *declarative* (or requirement) specification, which to some degree is indeterministic in nature. It is obvious that declarative specifications may be much more compact than the programs covered by them, even by orders of magnitude. We will come back to this point shortly.

A possibly suitable candidate among the various logical formalisms for legal purposes is some version of *description logic* (Bibel et al. 1993), such as *SHIQ* (Horrocks and Ulrike Sattler 2004), which has already been used extensively for formalizing e.g., medical knowledge. Using such a language, our task has two parts. One consists in building a terminology (i.e., of legal concepts) for describing the acts. The other consists in characterizing illegal acts using this terminology together with additional knowledge about, for example, relationships among different concepts. The terminology along with those relationships is also called an *ontology*. So what we need may also be described as an ontology along with the description (or declaration) of illegal acts.

The term “ontology” has a different connotation in jurisprudence, where it is used in its philosophical sense. Although there is of course some relationship to the underlying philosophical considerations (see Section 3.4 below), the main goal here is the formalization of a precisely defined, yet flexible net of the concepts on which law and legal reasoning are based.

Large ontologies have been built or are in the making (Mizoguchi 2004). One of the biggest is that underlying the CYC representation of common sense knowledge. CYC has more than a hundred thousand concepts in its ontology. But there are many more such large ontologies, e.g., Wordnet, Enterprise Ontology, Gene Ontology, Process Ontology, IEEE Standard Ontology, Cancer Ontology, and so forth. Legal ontologies have been developed (Gordon 2004b) but they are lagging behind in terms of size and completeness. Moreover, I am not aware of a legal ontology embedded in a larger knowledge system comparable to CYC.

There is no technical reason for this lag. Rather, the reason presumably lies in the fact that law is closer to the humanities and social sciences than to the natural sciences or engineering, which means that legal professionals are simply less aware of technical opportunities. Also, in contrast to domains like business and engineering, there is little or no economic incentive for change in the legal domain, due to the lack of a competitive market. With the global competition of regions and states this might be about to change, since the quality of the legal system is an important criterion for any prosperous area.

While I have mentioned description logics as one of the possible representational formalisms, it may still be debated which of the many alternatives

might be the most appropriate one for the legal domain. We will come back to these more technical issues in the subsequent section. The choice of language and formalism is important, but not to such an extent as to hinder the development of systems. Once the knowledge has been formally coded in some formalism, translation to another formalism may be performed more or less automatically.

Many authors have argued that a logical treatment of law is in principle unable to fully reflect the nature of legal argumentation. This is true if one has classical logic in its original form in mind. However, in the meantime all non-classical aspects of logic important for the law (and other applications), such as non-monotonicity, causality, time, vagueness, states of belief, probabilism, argumentation, pleading, theory formation, case-based reasoning, and several others, have been studied extensively. The maturity of the resulting formalisms is such that arguments of this kind are simply no longer appropriate. As with many other domains, Intellectics methodology is ready to be applied to the law.

Of course, formalizing a given piece of law potentially amounts to a lot of work if done by hand. Natural language understanding systems have reached a level of maturity which opens the perspective of supporting this formalization task to a substantial extent by machine. Similarly, linguistic techniques can be employed in building ontologies.

Another concern with automating law is its variability, since the law is in constant flux. So a given formalization must be kept synchronized with all these changes. By partially automating the acquisition of the formalized version from the natural language code, as just noted, this problem could eventually be overcome.

CONSEQUENCES OF A FORMALIZED LAW

So let us assume that a substantial part of some law has been coded the way we just indicated. What will have been gained? The advantages have been described many times and as early as two decades ago. What has changed since those early times is that most people now have already had first-hand experience of the advantages of coded information, for instance through the use of search engines (such as Google) which can give you correct answers to many of your questions, more or less immediately. So it might be worthwhile to repeat the arguments again now that more people might be more prepared to listen.

A first advantage of formalized law would be the possibility to use much more fine-tuned search engines to help judges and lawyers solve the problem, mentioned in the first section, of having to manually filter an over-abundance of data. That is, if both legislation and case law were precisely formalized,

search engines would be able to retrieve information relevant to a given case in a semantic way, avoiding the irrelevant data typically retrieved by purely syntactic search engines. This way cases, coded rules, theories, procedures, hierarchies of authority, norms and meta-rules would be accessible in a more directed way, based on systematically developed ontologies of concepts. What we are referring to here could be called the “legal semantic web.”

Further, if the law were formalized to a certain degree, it would become accessible to automatic consistency checks. Any man-made system has flaws. Early versions of large software systems suffer from tens of thousands of bugs. Unless one believes in miracles, the same is most likely true of larger bodies of laws and regulations. There are well-developed techniques for eliminating such inconsistencies. The legal domain, however, may require more elaborate methods than, for instance, those sufficient for first-order axiomatic mathematical systems.

A more basic, but also more technical issue concerns the optimization of the formalization. As is well familiar in jurisprudence, the choice of appropriate terms may make a great difference in the compactness and applicability of a law. Without the use of formalization, changing the choice of terms can take years or decades. Once part of the law has been formalized, existing mechanisms for terminology optimization can be applied to make the law much more compact, and this optimization could be achieved within days. The same applies to the rules in the legal code, once the terminology has been fixed; the set of rules could also be representationally optimized, using existing techniques. Note that Software Engineering (SE), which faces the same kind of problem, today has Computer-Aided Software Engineering (CASE) tools that support this kind of “refactoring” of specifications and, indirectly, computer code. Also note that optimization is superior to the kind of simplistic deregulation currently en vogue politically. Deregulation typically involves repealing existing statutes or regulations, without any substitution. While everyone applauds getting rid of the kind of truly ridiculous sections of code usually mentioned as examples, deregulation brings with it the danger of changing the meaning of legislation in undesirable ways.

A particularly important advantage of formalization has to do with the distinction we made between programs and their high-level declarative specifications. At the interface between a legal system and its users in the outside world, only the high-level declarative characterization of illegal acts is necessary and desirable. In contrast, it has become a particularly bad habit of legislatures to write into laws all kinds of little details at a low level of abstraction. This is one reason for the explosion in the number of laws, a situation deplored at the outset of this paper. In a formalization, the specification of the norm function could be radically optimized and simplified by phrasing the formalization as much as possible in a declarative and hierarchically structured way. In this way, laws would again become under-

standable to ordinary citizens. This again is analogous to ICT systems. Microsoft's Dynamic Systems Initiative, for example, envisions among other features a programming environment for formulating system definition models.

As we know from the ICT domain, declarative specifications are appropriate at the user interface but problematic in terms of efficiency. Therefore, in order to speed up the process of determining the norm value of some act (i.e., solving a legal case), internally the legal system could build-up program-like procedures satisfying the declarative specifications and link the two to each other. In addition, the declarative specifications could also be linked to the natural language legal code, resulting in what (Gordon 2004a) termed "Modelle von Rechtsquellen" (models of legal sources).

The main advantage of formalizing law is that legal cases can be solved with substantial support from an inference engine for the chosen formalism. Namely, solving a case logically means in a first approximation deducing its formal representation from the formalized law. (In more detail, abductive steps etc., are involved as well (Gordon 1995). See Section 3.4 below). The methods of automated deduction are so advanced that to some extent they should be feasible for supporting the practice of law. Note in this context that 80% of all legal disputes are settled "on the court house steps;" it is realistic to expect to be able to settle some of those "easy" cases in a more or less technology-driven way within the next two decades.

But technological support will not be restricted to the settlement or decision of cases. Other legal tasks requiring the analysis of facts and application of relevant law, including advocacy, pleading, advising, planning, legal drafting and public administration, could also benefit.

The consequences of such possibilities are so attractive that one can hardly understand why such little effort has been put into the development of this technology on a grander scale. For instance, imagine if, in a dispute with your neighbor, it would be better to consult your computer before heating up the discussion. You, or your lawyer, would describe your case and get the resulting value of the norm function right away, along with clear justifications or explanations as needed. Since your neighbor would do the same, there would be no point in continuing the dispute. Rather, it would make more sense to settle the case in accordance with the advice obtained using the system. I am not saying that we can have such systems for a variety of areas of law by tomorrow or next year. What I am saying is that *we have the formalisms and methods needed to realize a useful system of this kind within a midterm perspective.*

In fact, for limited domains such systems already exist and are in daily use, especially for tasks in public administration, such as filing tax declarations or applying for social benefits. In Australia, e.g., SoftLaw (<http://www.softlaw.com.au/>) has developed a system for veterans administration. Another example is an application currently under development in Germany for

Elternunterhaltszahlungsrecht (parent support law). The development and use of these applications has demonstrated that they are easy to build, can highly increase the productivity of public clerks, while improving the quality of their work, and make it much easier for citizens to access public services and take advantage of their rights. So these systems can save money and at the same time lead to better and more just results.

So we see that such technology can make the law more accessible to clerks in public agencies and especially to ordinary citizens. But in the midterm perspective it would also revolutionize and speed up the work of lawyers and judges dramatically. The quality of judgments and sentences would be greatly enhanced in all respects. Many of the problems listed at the outset of this paper would diminish to some extent. As we will discuss now in the next section, this includes problems with the legislative process (cf. Reinermann 2004).

TECHNOLOGICAL SUPPORT FOR LEGISLATIVE TASKS

The current practice of rulemaking consists of a process with three phases, labeled “notice,” “comments,” and “final rule.” As already mentioned in Section 2, legislation follows certain prescribed procedures consisting of a series of steps or hurdles that must be cleared. A functional conceptualization of this process is as revealing as the functional view of the primary role of the law given in that section. The report (Coglianese 2004, p. 30) structures the functions involved in rulemaking around the following tasks: idea, regulatory agenda, analysis and design, notice of proposed rule making (NPRM, e.g., in the Federal Register in the US) and comment, analysis and revision, rule publishing, enforcement and monitoring. The report also identifies relevant ICT for each function involved. Rather than discussing all these details, we want focus on the advantages of a declarative approach based on knowledge technology for selected subtasks.

The idea of democracy includes the rooting of the law in the thinking of the people rather than some elite, termed the *democratic legitimacy* of law. The integration of electronic comments into the regulatory process via a web page, as practiced e.g., by the US government, is a beginning (Coglianese 2004). Even here AI technology is already vital for coping with massive amounts of data, in the form of natural language-based search engines, and text categorization and summarization tools. Ideally, citizens would have to be polled for their opinions on proposed legislation. But one cannot expect everyone to understand abstract legal code. Alternatively, one could poll people for their opinions about particular legal cases and accumulate the results over longer periods of time, in order to abstract from the current zeitgeist. Using AI methods such as abduction or theory formation from

examples, the legal code could be generated so that it conforms with the results of such polls. At any rate, the regulatory process should become much more interactive and deliberative so as to increase its legitimacy.

Assuming the political will has been clarified in the way just described, the next step is to transform it into law. Rather than getting drowned in the details by “programming” legal code as is currently done, the results of the political discussion should be formalized declaratively at a level of abstraction which is accessible to ordinary people, such as elected representatives. At this declarative level, the integration of new laws into the existing legal codes, taking into consideration ethical and political values as well as expert knowledge from relevant fields, could be achieved more coherently. Also, such a declarative specification could leave open details so as to be able to flexibly adapt to contextual changes without affecting the basic code. This flexibility is especially important in view of the rapid progress of technology, where the law is lagging far behind and needs to confront new challenges (Miller 2005). Finally, a declarative specification of legislation would enable legal consequences to be simulated on computers, to detect undesired effects and improve the quality of legislation during the drafting phase, i.e., before the law is actually enacted. All these advantages contribute to making complex legal issues more manageable and accessible.

As we pointed out in Section 2, the legal system or, more generally, governance comprises much more than just the generation and application of legal code. There are many agents involved (see Gordon 2004a) and it features business processes like those of public administration. Therefore issues currently discussed under the heading of e-government are just as relevant for the legal domain as they are for the public sector in general. These include the optimization of administrative processes, keeping the public informed about new regulations, facilitating the understanding of rights and obligations in order to improve regulatory compliance, and many other issues, some of which we mentioned before.

Governments have actually taken notice of the importance of these issues, and are taking steps to improve the regulatory sector in this regard, mostly concerning the political and managerial, but rarely the legal dimensions. One of numerous examples falling into this category is Article 17 of the E-Commerce regulation of the EU, which aims to accelerate arbitration procedures by requiring them to be carried out electronically. Another is the E-Government Act in the US from the year 2002, which promotes the use of ICT by federal agencies for adjudicatory and rulemaking proceedings (Coglianese 2004). Finally, there is the E-Rulemaking Initiative of the US government, which considers e-rulemaking to be a promising area of innovation for regulatory agencies, with potentially important public benefits. As a consequence, the National Science Foundation (NSF) has launched a “digital government research program.” But all these measures more or less

ignore the fact that the concerned issues are only a part of the legal sector as a whole, and they all fail to address the core of the legal domain, namely the legal code.

A special kind of rulemaking, the drafting of contracts, is of particular relevance for business. One of the first successful prototype legal applications of Artificial Intelligence in Germany, KOKON, was built under my direction. KOKON demonstrated the feasibility of computer support for drafting contracts (Schneeberger 1992).

Another perspective concerns the elaboration of the facts underlying cases, i.e., the application of our fact-finding function. The function takes as input descriptions from various sources (evidence, the arguments of the plaintiff and defendant, precedent cases, contextual and commonsense knowledge, ethical principles, values, etc.) and constructs on the basis of this input a model of the facts of the case. The features of this construction process have been studied extensively in Intellectics. These include determining the common descriptive basis, as well as inconsistencies among the various sources, and the generalization and theory formation task of completing fragmental information pieces and bringing them into a coherent and realistic model of an act in the world and embedding this model of the case into the legal framework which is applicable to it. This generalization task also involves an optimization process, in order to construct a maximally coherent theory with a minimal number of speculative elements.

A short-term goal which could be tackled initially with Intellectics methodology is formulating a precise statement of the tasks underlying the fact-finding function, of the kind given for an analogous area (the tuning of metaheuristics) in (Birattari 2004, Ch. 3). While there are models for this function (Bench-Capon and Sartor 2003; Sartor 2005) in the literature, in contrast to this analogous area they lack an essential ingredient, namely a precise metric for theory coherence. Such a metric has been a requirement for any scientific enterprise since Galileo Galilei admonished: “*Misura ciò che è misurabile, e rendi misurabile ciò che non lo è*” (measure the measurable and render measurable that which is not yet so). The probabilistic features elaborated in machine learning may form the basis for such a metric in law, just as they do in metaheuristics.

Intellectics could further clarify the role of peculiarities with the human way of solving this function, which is currently based mostly on commonsense psychology, rather than scientific facts from Cognitive Science. And, in the long-term, Intellectics could contribute systems supporting daily legal work. Thereby, a clear line must be drawn between the clarifying support provided by such systems and the final interpretation and decision of legal cases, which must be left to judges in all but trivial cases.

Yet another perspective concerns law enforcement applications of “ambient intelligence” technology, now a focus of research and development

in Europe. (It is no coincidence that ambient intelligence shares the abbreviation “AI” with artificial intelligence, as they have much in common.) Namely, if computation becomes ubiquitous, as envisioned by ambient intelligence, illegal behavior, such as business fraud, could be detected more easily, resulting in more justice and security. The issues involved include privacy, surveillance, the “USA PATRIOT Act” and others. Moreover, we may anticipate new types of legal penalties, based on a wider spectrum of freedom deprivation methods and fine-tuned degrees of coercion.

A long-term perspective for legal knowledge systems is supporting the building of “A New World Order” (Slaughter 2004) in which international networks of public agencies, including courts, closely cooperate, leading to global governance while respecting differing national laws and values. “*The greater use of foreign material affords another source, another tool for the construction of better judgments. ... The greater use of foreign materials by courts and counsel in all countries can, I think, only enhance their effectiveness and sophistication.*” said Justice G.V. La Forest of the Canadian Supreme Court. This thought has also been expressed, in different words, by US Chief Justice William Rehnquist.

The trend towards globalization, also in the realm of law, has to a large extent been triggered by technological developments, not least from those in AI. Today, globalization in the law is facilitated by IC technology, for example by the data-bases of Lexis-Nexis and Westlaw or websites such as CODICES, which publish legislation and court decisions from around the world. The technological perspective I have presented here has the potential to support the semantic filtering of the masses of legal material, and facilitate the learning, understanding and appreciation of the multitude of rules (and values – see below) in different countries. This is indispensable for finding commonalities and tolerating differences, whereas ignorance is fertile soil for hatred and aggression.

Harmonization and convergence by “persuasive authority” hopefully will be the result of this process of an emerging global jurisprudence. Such an outcome is especially important in the European context, with its 25 national legal systems in addition to the European system on top of them. Note that in the EU each European law must be implemented as national law in each member state, and thus has to be “translated” into 25 different legal codes. The situation in the US is perhaps even more complex, since the laws of the individual states coexist with federal law, resulting in many inconsistencies.

Another long-term perspective is Intellectics support for ethical norms and societal values, analogous to its support for legal norms. Values could be modeled in a hierarchical way, from the most abstract concepts of freedom, justice, or peace down to more concrete values, such as the preservation of natural resources or sustainable development. Especially in the current atmosphere of fear due to terrorist threats, politicians are quick to call for a

return to traditional values without specifying what they actually mean. Determining these values should, again, be a matter for all the people and not just some elite, even if democratically elected.

A GLIMPSE AT THE PHILOSOPHY OF LEGAL REASONING

We already noted that the legal community has been slow in grasping the potential of Intellectics technology and we mentioned ignorance and lack of competition as possible reasons. Adopted reflexes to misunderstood terms may be another factor in this complex set of social attitudes. We briefly mention in this section a couple of factors of this kind and refer the reader to Chapter 3 of Gordon (1995) for a more detailed analysis as well as for references concerning the subsequent discussion.

One misunderstanding concerns the term “ontology,” as already noted. The semantics of this term as it is used in Intellectics inherits some of its original philosophical meaning, but lays more emphasis on the structural properties of nets of concepts than on the (fuzzy) philosophical foundations of the basic concepts themselves. Those in the legal community who understand this difference may nonetheless fall into the trap of misleadingly identifying the use of ontologies with *conceptualism* (Begriffsjurisprudenz), the dominant legal philosophy in Germany of the second half of the 19th century, which usually is ridiculed today. But, as we have explained, an ontology in the Intellectics sense is neither fixed nor can it provide a complete account of legal reasoning. Rather, an ontology is just the terminological basis of any more comprehensive formalization, covering many further features of legal reasoning, as discussed above.

Another misunderstanding concerns deduction. The first attempts to model legal reasoning in the middle of the past century viewed legal rules as a body of axioms and considered the legal analysis of a case to be a deductive derivation from these axioms. Only through further investigations did it become clear that this model is too poor to capture the essence of legal reasoning. In Section 3.1 we mentioned many non-classical features of legal reasoning, such as the incompleteness and vagueness of concept definitions (open-texture, “Porösität der Begriffe”), the need for non-monotonicity, abduction, theory formation etc. However, we now know how to cope with these features as well and, in fact, it has turned out that one can still model them in a deductive way, at the meta-level. So the original, deductive view of legal reasoning was not as ridiculous as some might let it appear. In other words, deduction remains at the heart of models of legal reasoning.

A final aspect to be mentioned here is the interpretation of formal text. Any formalisation of the real world creates a virtual world different from the real one, as does any description in natural language. This virtual world

necessarily can be no more than an approximation of the real one, which again is true for the virtual world created by natural language code. Whether the natural language approximation is the better of the two, as is claimed in hermeneutics, remains an open question. We cannot know the answer until experiments with formal virtual worlds are carried out in domains like the law. All prior experiments of this kind in other areas (e.g., in medicine, mathematics, and chemistry) have not provided any evidence in favor of the hermeneutical position.

Short-term recommendations

The vision of a technology-enhanced legal system remains, as demonstrated in the previous section, and we know better than ever how to realize it. Furthermore, with the dramatic advancement of ICT the chances for realizing this vision are improving by the day. And in fact there are already a number of very helpful systems in daily use.

In the short-term, one should concentrate especially on maximizing visibility through the building of success systems. This could be achieved, for instance, with systems supporting the interaction of citizens with public administration, in areas such as tax law, traffic law, or social benefits law). Such systems will help empower citizens to exercise their rights, accelerate administrative procedures and improve the quality of interaction in various other ways, while saving more money than they cost. In our times, when public budgets are tight and the legal workload is rising, and hence, placing an increasing burden on the budget, particularly these economic benefits should be taken seriously.

But our vision goes beyond such restricted systems. An important step toward the long-term vision would be the building of a legal ontology and the standardization of the way the legal knowledge is coded in such a system. This way, components could be shared rather than reimplemented unnecessarily. In the automated deduction community, a similar standard had a tremendous impact on subsequent progress. Standardization may be seen as the beginning of a more comprehensive system, in the midterm, partially achieved by aggregating existing components.

This is not to say there are no obstacles on the road toward the realization of more general systems. First of all, representing the entire law in some formal system is an “AI-complete” problem, in the sense that this would require achieving true artificial intelligence, something which certainly will not happen in the near future. But even partial steps towards systems with a limited scope and intelligence would already be of great help, as many partially intelligent systems in use demonstrate. “The” AI problem is not one to be solved in an “all

or nothing” fashion, as is sometimes maintained in the popular press, but rather will be achieved incrementally, with ever better solutions.

The second major obstacle is that these more general systems become useful only after some critical level of performance and complexity has been reached. In other words, a challenging task must be successfully completed before society and the legal profession will be able to see any convincing fruits beyond the specialized systems already in use. Only with public funding can this hurdle be overcome.

Thirdly, it must of course be admitted that there are still a number of research questions with unresolved details. In the remainder of this paper I would like to mention some of these questions (see also Sartor 2005).

For instance, let us come back to the issue of choosing a formalism for the representation of legal code. It can be selected from numerous alternative formalisms. One of the points of consideration in this choice is the fact that the definitions of legal terms often include transitional statements. For instance, in the German civil law, the *Bürgerliches Gesetzbuch* (Schumann 1980), § 1 reads: “Die Rechtsfähigkeit des Menschen beginnt mit der Vollendung der Geburt” (the legal status of a human begins with the completion of birth). So we are talking here about a concept which is limited in its validity to a certain time period, beginning at the point where the “act” of the mother delivering the baby successfully ends. Description logics do not readily suggest a way how to represent such transitions and I have not seen an answer to this particular question in the respective literature.

However, the problem as such has extensively been studied in the context of other formalisms. For instance, the huge CYC system has of course successfully dealt with this issue. The oldest formalism specifically addressing transitions is the situation calculus, whose main idea of adding a situational parameter could easily be carried over to description logic. But for several reasons this formalism has not been used much in practical applications. Another popular formalism for the same purpose is logic programming, e.g., (disjunctive) logic programs with stable semantics (Gelfond and Lifschitz 1988), again enriched by situational parameters. Some more recent work, such as (Bibel 2004; Bochman 2004), favors a causal approach. Here transitions are coded as separate rules, in addition to the classical logical relationships. The advantage of this causal approach is that they cover many aspects of knowledge representation at the same time, including transitions and causality, non-monotonicity and imprecision, and time.

These considerations show that, as a first step towards the representation of law, one has to test several of the existing formalisms for their adequacy for representing law in a suitable way. Once the choice for some formalism has been made, it is to be expected that some further work will be necessary to figure out for more intricate legal clauses how to represent them in the formalism. This may occasionally require extensions borrowed from other

formalisms. So although building a legal knowledge system will require substantial efforts, this effort will be worthwhile, considering the expected benefits. But without pressure from actual applications of these formalisms in the legal domain, the necessary integration of features from different formalisms will be achieved much more slowly.

The chances are good that research and development will be supported with public funds. For instance, during the preparation of the next call for the 6th Framework Programme of the European Commission, as well as the design of the 7th Framework Programme, a report (Bibel et al. 2004) was written which includes a short perspective of future legal systems (Section 2.3.4) and, among its recommendations, explicitly mentions legal knowledge bases (Sections 0.2.4.2 and 5.2.4.2). In the recent 4th Call of the European IST Program, we already find relevant issues, e.g., in “Strategic Objective 2.4.9: ICT research for innovative Government.” So it is now in the hands of the research community to seize this opportunity, by submitting proposals for larger projects in the legal area. Of course this is not the place to address this issue in further detail.

Conclusions

In this paper, we have discussed the problem of the increasing complexity of law and outlined an Intellectics approach to partially coping with these problems. Our focus was on a declarative approach within the legal dimension of governance, centered around legal knowledge systems. This approach is meant to be complementary to current efforts under the banner of e-governance, e-rulemaking, etc.

This ICT approach has the potential to achieve higher quality rules, make the law accessible to ordinary citizens, not only the wealthy and educated, improve the level of compliance, and to foster greater and deeper public participation, thereby improving the transparency of the rulemaking process. As current systems (such as those by SoftLaw) demonstrate, legal knowledge systems are relatively easy to build and can lead to high productivity increases in public administration, to an improved quality of service, and to much easier access to these services by citizens, allowing them to more fully take advantage of their rights. So altogether legal knowledge systems can save money and, at the same time, lead to better results and more justice.

As exemplified by the extremely useful systems already in existence, Intellectics methods and technology have reached a level of maturity which allows their immediate application in domains such as the law. Yet further progress requires challenging research issues to be addressed. One of these issues is the integration of various features provided so far only in separate formalisms. The best way to accelerate this integration process would be to generate demand by developing applications in domains such as the law. Not

only would this research greatly benefit the legal field, but also Intellectics would learn from such challenging applications.

Since such applications requires a fresh view of the legal system as a whole, progress also depends on research contributions from legal theory and the social sciences, using both formal and empirical methods, as has been argued in an EU report (Bibel et al. 2004) in a more general context.

In terms of my grander vision, i.e., far beyond the limitations of current systems, we are talking, of course, of major research and development efforts requiring larger amounts of funding. However, if the necessary funding is compared to the amounts made available for huge projects in some other fields, such as plasma or particle physics, they appear relatively modest, especially considering the anticipated midterm benefits for society.

“*Vor Gericht und auf hoher See ist der Mensch in Gottes Hand*” (in court and at sea man is in God’s hand) is a popular saying. Only the second half of this saying is still true, since boat trips of any kind have become rather safe thanks to modern technology. It is high time to address the first half of this proverb, especially now that we have the technology for substantially improving the situation.

Acknowledgements

I am grateful to a number of people for their help in the preparation of the material in this paper and for their comments, including Tom Gordon, Raoul Kneucker, Friedrich Lachmayer, Sonia Miller, Filippo Ranieri, Helmut Rübmann and Joachim Wenzel. Especially Tom Gordon has made substantial contributions, first of all by inviting me for a presentation at Jurix 2004 which laid the basis for the paper, but also by numerous interactions on the issues discussed here, and not least by polishing numerous of my weak English phrases. I am deeply indebted to him, while, of course, all possibly remaining flaws in the text are my responsibility. I also thank the organizers of Jurix 2004 for the invitation and editor Giovanni Sartor for his interest in the material and for his prompt handling of the submission.

References

- Bibel, W., Andler, D., da Costa, O., Küppers, G. and Pearson, I. D. (2004). *Converging Technologies and the Natural, Social and Cultural World*. Technical Report, Commission of the EU. http://europa.eu.int/comm/research/conferences/2004/ntw/pdf/sig4_en.pdf.
- Bench-Capon, T. and Sartor, G. (2003). A Model of Legal Reasoning with Cases Incorporating Theories and Values. *Artificial Intelligence Journal*, 150(1–2): 97–143.
- Bibel, W., Hölldobler, S. and Schaub, T. (1993). *Wissensrepräsentation und Inferenz*. Vieweg Verlag: Braunschweig.
- Bibel, W. (2003). *Lehren vom Leben – Essays über Mensch und Gesellschaft*. Sozialwissenschaft. Deutscher Universitäts-Verlag: Wiesbaden.

- Bibel, W. (2004). Transition Logic Revisited. (Submitted).
- Birattari, M. (2004). The Problem of Tuning Metaheuristics as Seen from a Machine Learning Perspective. Université Libre de Bruxelles: Brussels. PhD Thesis.
- Bochman, A. (2004). A Causal Approach to Non-Monotonic Reasoning. *Artificial Intelligence Journal*, 160: 105–143.
- Coglianesi, C. (2004). E-Rulemaking: Information Technology and Regulatory Policy – New Directions in Digital Government Research. Regulatory Policy Program Report No. RPP-05, Harvard University, John F. Kennedy School of Government, Center for Business and Government: Cambridge MA.
- Eckhoff, T. and Kristian Sundby, N. (1988). Rechtssysteme – Eine systemtheoretische Einführung in die Rechtstheorie. Duncker und Humblot: Berlin.
- Gelfond, M. and Lifschitz, V. (1988). The Stable Model Semantics for Logic Programming: In Bowen, K. and Kowalski, R. (eds) Fifth International Conference and Symposium on Logic Programming., Cambridge MA, S. 1070–1080. MIT Press.
- Gordon, T. F. (1995). The Pleadings Game – An Artificial Intelligence Model of Procedural Justice. Kluwer Academic Publishers: Dordrecht.
- Gordon, T. F. (2004a). Die Bedeutung von eGovernance für die Öffentliche Verwaltung. *Verwaltung und Management*, 10(5): 258–263.
- Gordon, T. F. (ed.) (2004b). Legal Knowledge and Information Systems – Jurix 2004: The Seventeenth Annual Conference, Frontiers in Artificial Intelligence and Applications, Amsterdam. IOS Press: Jurix.
- Horrocks, I. R. and Sattler, U. (2004). Decidability of *SHIQ* with Complex Inclusion Axioms. *Artificial Intelligence Journal*, 160: 79–104.
- MacCormick, N. (1993). Beyond the Sovereign State. *Modern Law Review* 56.
- Miller, S. E. (2005). Science Confronts the Law: In Roco, M. C. and Bainbridge, W. S. (eds) NBIC Convergence. Kluwer, Dordrecht NL.
- Mizoguchi, R. (2004). Tutorial on Ontological Engineering. Part 2: Ontology Development, Tools and Languages. *New Generation Computing*, 22(1): 61–96.
- Rissland, E. L., Ashley, K. D. and Loui, R. P. (2003). AI and Law: A Fruitful Synergy. *Artificial Intelligence Journal*, 150: 1–15.
- Reinermann, H. (2004). Editorial. *Verwaltung und Management* (2).
- Russell, S. J. and Norvig, P. (2003). *Artificial Intelligence: A Modern Approach*. (2nd ed.). Prentice Hall Series in Artificial Intelligence. Pearson Education: Upper Saddle River, NJ.
- Sartor, G. (2005). *Legal Reasoning: A Cognitive Approach to the Law Treatise*. Springer: New York.
- Schumann, E. (ed.) (1980). *Bürgerliches Gesetzbuch – BGB*. Goldmann Verlag: München.
- Schneeberger, J. (1992). Kokon: In Shapiro, S. C. (eds) *Encyclopedia of Artificial Intelligence*, Volume 1, S. 759. 2nd edition John Wiley, New York NY.
- Slaughter, A. -M. (2004). *A New World Order*. Princeton University Press: Princeton.