

Using Ontologies for Comparing and Harmonizing Legislation

Alexander Boer
Department of Computer
Science & Law
University of Amsterdam
Amsterdam, The Netherlands
aboer@lri.jur.uva.nl

Tom van Engers
Dutch Tax and Customs
Administration
Utrecht, The Netherlands
t.m.van.engers@acm.org

Radboud Winkels
Department of Computer
Science & Law
University of Amsterdam
Amsterdam, The Netherlands
aboer@lri.jur.uva.nl

ABSTRACT

In the last decades the interest in the problem of comparing and harmonizing legislation has been steadily increasing. One reason is the increasing legal convergence between governments in the European Union, and the increasing traffic of people over borders of jurisdictions. Another reason is the increasing globalization of companies; Products and services are offered in many jurisdictions at the same time, and the product or service has to meet the provisions of all jurisdictions in which it is offered. In the E-POWER project relevant tax legislation and business processes are modeled in UML to improve the speed and efficiency with which the Dutch Tax and Customs Administration can implement decision support systems for internal use and for its clients. These conceptual models have also proven their usefulness for efficient and effective analysis of draft legislation. We are currently researching whether conceptual modeling can also be used to compare ‘similar’ legislation from different jurisdictions to improve the capacity of the Dutch Tax and Customs Administration to react to future consequences of increased movement of people, products, and money between EU member states and increased harmonization between tax authorities in Europe. In addition, addressing the problem of comparing models is also expected to improve our methodology for modeling legislation. This paper discusses problems and requirements of comparing legislation as we understand them now, and attempts to relate them to relevant research.

1. INTRODUCTION

In the last decades the interest in the problem of comparing and harmonizing legislation has been steadily increasing. There are a number of reasons for this development. One reason is the increasing legal convergence between governments in the European Union, and the increasing traffic of people over borders of jurisdictions. This leads to an increased need of administrations to know and understand legislation of friendly governments to be able to assist citizens and reduce negative consequences of movement. Another reason is the increasing globalization of companies; Products and services are offered in many jurisdictions at the same time, and

the product or service has to meet the provisions of all jurisdictions in which it is offered. Different regulations can also lead to differences in competitiveness for the product or service. For a financial product, for instance, it is considered important to qualify for tax deductions that make the product more attractive. For a medical product it is important to know whether it can be sold over the counter without a prescription.

The increase in comparing legal systems is evidenced by the number of consultancy firms that advertise their knowledge of multiple legal systems to companies. In addition, there are a number of initiatives for constructing international legal ontologies that expose the subsumption relations between legal vocabulary in multiple jurisdictions (e.g. [16]). It has also been suggested in recent literature [11] that the problem of comparing legislation is, at least if it concerns versions of a regulation in time in the same legal system, a special case of the general problem of ontology integration [6]. In this paper we attempt to show that that this is a promising approach, but that the problem of ‘comparing’ legislation is broader than the problem of ontology integration.

In the E-POWER project we model relevant tax legislation and business processes in Unified Modeling Language (UML) to improve the speed and efficiency with which the Dutch Tax and Customs Administration (DTCA) can implement decision support systems for internal use and for its clients. The ontologies have also proven their usefulness for efficient and effective analysis of draft legislation, allowing the DTCA to give immediate feedback to drafts of the new income tax law of 2001 [20].

We are currently researching whether the available ontologies can also be used to compare ‘similar’ legislation from different jurisdictions. Employees in the DTCA are increasingly confronted with requests that require them to understand European regulations and directives, and regulations of other EU member states. Better insight in the process of modeling and comparing legislation from different legislators is expected to improve the capacity of the DTCA to react to future consequences of increased movement of people, products, and money between EU member states and increased harmonization between tax authorities in Europe. In addition, the discovery of the requirements of comparing models is also expected to result in a more principled, more robust, and language-independent methodology for modeling legislation, because it decouples the model of legislation from the purpose for which it was modeled and makes assumptions about the surrounding context, and knowledge expected from users, explicit. Modeling one’s own regulations of-

ten fails to achieve this level of context-independence.

Comparing similar regulations from multiple jurisdictions is of course not the same as comparing the legal systems to which the documents belong. Two countries may for instance contain almost a copy of the same crime description in their respective penal codes, but it is also important to know whether both are backed up by similarly likely sanctions. Other important differences are the degree of civil servant discretion, corruption, and different constitutional arrangements influencing when and how regulations are applied. In civil and common law systems, for instance, regulations are often applied differently and are therefore written in completely different styles. Different jurisdictions, cultures, and languages, and the philosophical problems of comparing legal systems are the main subject of the field of Comparative Law. In this field of research, directly comparing two regulations is a very uncommon approach because everyone is acutely aware of the philosophical problems. Although the E-POWER project limits itself strictly to comparing legislation, and not legal systems, we are aware that the two cannot be clearly separated. The concepts used in legislation reflect the legal system that drafted and uses that legislation. To find similarities and differences between concepts one needs to know to what things the concepts refer, and to compare legislation one needs to know its real effects.

We believe it is realistic to compare regulations of very similar legal systems in continental Europe, certainly in areas like taxation where uniformity of approach has had a clear and obvious economic value for a long time. In other cases the conceptual differences may be so huge that no clear correspondence between legislation can be established, and the effort required for modeling the surrounding legal systems outweighs the benefits of comparing legislation in a formal way. This paper discusses the problems we are aware of, and the requirements for comparing legislation we have formulated. The next sections discuss integration of legal ontologies, the problem of establishing ‘similarity’ between regulations and concepts, the concepts needed to describe and explain legislation itself, and the specific problems related to comparing legal ontologies of limited scope. Some remaining problems with our proposed method for comparing ontologies of legislation will be discussed in the last section.

1.1 Integration of Legal Ontologies

A number of international efforts to establish standards for legal XML have recognized the problem of mapping legal vocabularies to each other and proposed standard ontologies of international legal vocabulary. Because XML has become the de facto international syntax standard for legal document exchange between organisations, and many organisations have plans for offering web services to international customers, there is an increasing recognition of the need for translation of legal jargon. The European LEXML consortium¹[16] has initiated the LEXML RDF Dictionary project, and the OASIS LegalXML working group² has a committee for a very similar dictionary project. The difference between priorities set in the LegalXML approach, which is primarily rooted in the common law jurisdiction of the U.S., and the LEXML approach, rooted in the civil law jurisdictions of several continental northwestern European countries, illustrates the barrier between civil law and common law in this field. The METALex consortium

¹<http://www.lexml.de>

²<http://www.legalxml.com>

[5]³, also European and associated with LEXML, has a similar ontology subproject. The LEXML and LegalXML dictionaries are intended to facilitate translation between existing legally relevant XML data structures defined in a DTD or XML Schema describing the structure of a particular kind of legal document, e.g. a verdict or contract. The METALex initiative tries to set generic and simple jurisdiction-independent XML standards for legal documents, and allows jurisdiction-specific extensions to its XML schemas that can be translated to the generic vocabulary. The ontology is intended for translation of jurisdiction-independent XML documents to a normative common jurisdiction-independent vocabulary that can be explained to all. The E-POWER modeling and annotation tools write METALex Regulation XML.

These three projects all start from the assumption that a taxonomy of legal concepts is the obvious way to explain the relations between concepts from multiple jurisdictions. The problem of comparing legal systems is thus reduced to an ontology integration problem; Multiple jurisdiction-specific ontologies must be integrated into one international ontology. Approaches and tools for integration of ontologies are listed and categorized in depth in [18]. The LEXML philosophy assumes that it is possible to use a bottom-up approach, where the structure of the international ontology is discovered along the way. METALex starts from the assumption that the integration problem requires a global ontology of top categories that distinguishes concepts by formal criteria such as rigidity, existential dependence, etc. to provide a conceptual backbone guiding the insertion of natural and underspecified concepts. The Framework for Ontology Integration [6] exemplifies this approach that is more common in the field of ontology integration.

This approach of integrating ontologies representing two regulations has been tried recently [11] to model norm dynamics. In [11] norms in a regulation are represented as standard concepts that derive their meaning from other concepts in that same regulation. The regulation is thus represented as a ‘local’ ontology containing concept definitions and the problem of integrating two ‘local’ ontologies representing two different regulations is reduced to a problem of ontology integration. They observe that integrating the ‘local’ ontologies representing two versions of a regulation in time is a way to compare and find differences, and that ontology integration is less complex if both ‘local’ ontologies to be integrated have been built with reference to a shared ‘global ontology’. The problem of listing the changes between the two versions is then reduced to a question of discovering changed subsumption relationships between concepts. Generic comparison relationships based on subsumption between concepts in a description classifier are specialized to comparison relationships between two norms, and the description classifier is used to decide whether norms are retained, amended, derogated, or extended. In [11] the problem of comparing legislation is reduced to a special case of the general problem of ontology integration, inspired by other work [10] on the DOLCE ontology, which is an example of a global ontology.

In [11] the scope of the problem is limited to norm dynamics, and different jurisdictions are not taken into account. We agree with the viewpoint on integration of norm systems in [11], and the nature of norms. ‘Norm’ is a concept, and norms in a regulation are instances of the concept ‘Norm’, but norms are at the same time concepts themselves (e.g. murder, theft) even though they are very often unnamed. [11] do not account for norm conflict resolution

³<http://www.metalex.nl>; The authors contribute to this consortium.

strategies for establishing the validity of norms. We sketch such a component in this paper. In addition, we claim that the problem of ‘comparing’ legislation as it occurs in practice is broader than the problem of ontology integration. The ontology integration view does not account for the purpose of the comparison in a satisfactory way, and that explains why publications in the field of Comparative Law do not seem to be related to ontology integration at all.

2. COMPARING LEGISLATION

Similar regulations are compared for a number of different purposes, and in most cases the comparison is not at all that complex because the surrounding legal system remains the same. Regulations are for instance compared for:

Policy Comparison Proposals for a regulation addressing the same problem are compared to judge which one is better according to preconceived norms of analysis.

Forecasting and Reconstruction Two versions of the same regulation in time are compared to determine the effects (costs and benefits) of changes of legislation on behaviour, products, etc.

Migration Two regulations addressing ‘similar’ things in different jurisdictions are compared to inform others about the effects (costs and benefits) of moving themselves, their property, products, or services over the borders of a jurisdiction.

Harmonization efforts combine these tasks; It often aims to minimize the costs of migration, distinguishes the good from the bad legislation to repair the latter, and forecasts the costs caused by the changes it proposes in order to be able to minimize them. To harmonize legislation one has to quantify and prioritize costs and benefits for stakeholders with a variety of norms of analysis.

It is clear that comparing legislation is not the same as comparing two ‘local’ ontologies of legislation because legislation is usually compared with implicit norms in mind. Harmonization is the most complex context for comparisons because the regulated worlds are not the same, and the mapping of legal and commonsense concepts over the borders of a jurisdiction is not as straightforward as the ontology integration perspective applied to harmonization would suggest.

2.1 Similarity of Legal Concepts

Obviously, when two different legal cultures and systems come into contact, there is significant potential for misunderstanding. Translation of legal documents to a foreign language illustrates the nature of the problem. If legal texts are for instance conceived and drafted in Dutch, based on the concepts of Dutch law, and then translated into English, the result is an English text from a structural and linguistic point of view, but the text is semantically rooted in the Dutch law and society. An explanatory document on Dutch law that was directly drafted in English is also still rooted in Dutch legal concepts. This problem is not unique for translation of *legal* texts, but it is especially acute in this context. In fact, the ‘translation’ problem between legal systems also exists for U.S., E.U., and British legal English, and communication between laity and professional in general. The translation must strike a balance between using the concepts the audience already knows, and teaching the definitions of unfamiliar legal concepts.

To translate from legal jargon to ‘commonsense’ vocabulary, or Dutch to English, one needs to explain the concepts from one vocabulary in terms of the other. Translators use three general strategies to explain concepts and institutions bound to a particular legal system:

Literal translation Take a dictionary, break a composite word into its constituent parts if necessary, and translate word for word. For instance, the Dutch term ‘bestuursorgaan’ becomes, according to the dictionary, either ‘government organ’, ‘administrative organ’, ‘government body’, or ‘administrative body’.

Transfer to a similar concept The Dutch concept ‘bestuursorgaan’ comes from general administrative law (a literal translation), which is applicable to all institutions labelled ‘bestuursorgaan’. It regulates i.a. procedural requirements for administrative decisionmaking, appeals against administrative decisions, and delegation of decisionmaking competence⁴. The concept ‘public body’ is used in a very similar sense in the United Kingdom and the English vocabulary of the European Union. A translator acquainted with these legal systems may therefore substitute ‘public body’ for ‘bestuursorgaan’.

Periphrasis Explaining an unfamiliar concept by a defining gloss using other concepts. Observing that translating ‘bestuursorgaan’ with ‘public body’ fails to convey the limits of the concept, a periphrasis (or circumlocution) may be more suitable: *bestuursorgaan* – an administrative body of a public legal person, including natural persons exercising public authority, bodies or colleges without legal personality, and bodies of a private legal person exercising public authority insofar as they as they are making decisions based on a competence attributed by law, excluding the assembly of the States-General, assembly of a house of the States-General, the Judicial Organisation, the Council of State, the Court of Auditors, the National Ombudsman, etc.

Indiscriminate application of periphrasis results more verbose text that lacks a semantic grounding in the vocabulary of the target audience. Indiscriminate use of transfer, on the other hand, results in a text that is misleading from an information perspective. The transfer may be more attractive to the target readers, but the readers are not aware of omissions. The most obvious procedure is thus to apply periphrasis until unambiguous transfer becomes possible.

Whether transfer is unambiguous depends on the knowledge attributed to the intended reader and the problem at hand. If you are thinking of preparing a complaint against a public body in the Netherlands, an ontology that classifies ‘public body’ and ‘bestuursorgaan’ as similar has little value.

Intuitive ‘similarity’ is for purposes of Comparative Law classified by [25] as *extensional*, *immanent*, and *functional*. Concepts are the same in an *extensional* sense if their definitions are logically equivalent. Suppose that statements that harm someone’s honour or dignity constitute insult, while untrue statements that harm someone’s honour or dignity constitute defamation. Insult and defamation as described here are *immanent* concepts, abstracted from particular legal systems and the ‘extensional’ intersection (overlap or ‘common ground’) of specific definitions found there. In the *extensional*

⁴In E.U. English. Others may prefer ‘power’

sense, the definition of ‘insult’ subsumes the definition of ‘defamation’; All defamations are by definition insults, but not the other way around.

Some western countries provide a civil remedy for defamation, while other countries (including the Netherlands) may treat the very same case as a criminal insult, arguing that establishing the truth criterium often leads to extra harm for the victim and is usually irrelevant because ‘acting in the public interest’ is a valid defense for both. Although civil defamation and criminal insult are different even with regard to the nature of the procedures to be followed and the legal consequences, they are *functionally* similar because they address a ‘similar’ problem in a society. In this sense concepts can be functionally ‘similar’ from the point of view of Comparative Law, even though the definitions are completely different.

2.2 Similar Norm Systems

Because functional similarity must be based on either the perceived intentions of the legislator, or the perceived causal effect of the concept on the ‘similar problem’ of different societies, this notion of similarity is not very useful for automatic reasoning. For Knowledge Engineering purposes, we want to separate extensional and immanent views of concepts from values and effects. The problem is, however, that comparative analyses made by scholars in Comparative Law that reveal ‘interesting’ differences between laws are usually based on functional ‘similarity’. To create useful computerized support, the implicit underlying values must be made explicit.

Two variations of functional similarity can be distinguished:

Effect on behaviour Two norm systems can be considered similar if they cause similar behaviour of the group of agents addressed by the respective norm systems.

Intended effect on behaviour Two norm systems can be considered similar if they are both intended by the legislator to cause similar behaviour of the group of agents addressed by the respective norm systems.

The distinction between both becomes blurry if you make assumptions about the effect of norm systems on behaviour without validating them (confusing intentions with effects). That the two are different can be illustrated with a diesel fuel tax in two different countries A and B. Country A has a diesel fuel consumption tax of 50%, and country B has a diesel fuel consumption tax of 10% and an environmental diesel fuel tax of 40%. The intentions of the taxes are different. The consumption tax is intended to generate income, the environmental tax to encourage evasion behaviour by consumers of diesel fuel and maybe as a bill for health problems (assuming health is publicly funded) and pollution caused by diesel fuel. Empirical research may support the hypothesis that both taxes cause the exact same evasion behaviour and generate the same income. From an economic viewpoint they are then ‘similar’. At the same time they may be different because country B may have committed itself to spending the generated income in the environmental regulation to specific policies.

Our thesis is that the similarity that we seek to establish between regulations should be interpreted as being similarity in the changes they cause, or are intended to cause, in behaviour.

3. THE ROLE OF LEGISLATION

To explain how regulations affect the real world, and why one regulation is better than the other, we need concepts that explain the role that the text of the regulation plays in a jurisdiction. One example of a generic ontology of the content of legal sources for the design of legal expert systems is the *Functional Ontology of Law* of Valente (cf. [23, 22]). This ontology was also one of the first to stress the importance of semantically grounding legal concepts in shared commonsense concepts.

The key concept for understanding legislation is the *norm*; If you are describing two people playing what appears to be a chess game, for instance, and one player moves a pawn backwards, you may infer that it is not a chess game after all; The move contradicts the hypothesis that it is a chess game. If you are the other player, however, and you believe you are playing a chess game, then you will consider the move illegal – a violation of a norm. In this case you view the rules of chess as a prescription of what ought to happen during the game. The norm prescribes what ought to happen. Norms only regulate human behaviour; The eruption of a volcano may very well be undesirable, but it makes no sense to try to prescribe how it ought to behave. Norms are intended to change people’s preferences between choices, to interfere with basic economic behaviour by changing “the rules” of the game. Norms therefore presuppose an understanding of what the norm entails and legal systems only function if the norms are known.

Legal norms are therefore made explicit in a document (contract, regulation, jurisprudence, etc). The document prescribes behaviour to agents assigned a certain role, for instance the owner and user of a road, or posits norms for the creation of artefacts – things made by agents, for example ships, works of art, and tax forms. In other cases it constrains or defines procedures for actions or transactions by agents, for example a survey, hearing, or purchase. In each case, the norm (indirectly) makes a statement about human choice, it imposes a preference. The situations constrained in the norms expressed in legislation are described in terms of a mixture of well-understood ‘commonsense’ concepts and the legal concepts discussed in the previous section, that are eventually defined in terms of ‘commonsense’ concepts. Norms can restrict the behaviour of everyone; The application of codified norms by a judge to others, and the imposition of sanctions⁵, is also regulated by norms.

Norms occur in three main variants in natural language: obligations, permissions, and prohibitions. The use of verbs is not indicative of whether something is a norm or not. Both the legislator and the sociologist can utter, for example, the deontic sentence “Tanks adjacent to the hull are not used to store fuel oil”, and in neither case the sentence would be semantically ambiguous. In the first case, the legislator constitutes a rule, assuming as thesis a tank adjacent to the hull, prescribing how it should be used. In the second case, the sociologist verifies a regularity, analyzes the relation between the norm and the social reality describing a situation.

Deontic verbs are usually treated as separate modalities in deontic logics, but another way of formalizing them treats them as expressions of preferences of the legislator imposing a partial ordering with just two values – allowed and disallowed – on situations (see for instance [24]). If one models the norm as a preference relation imposing a partial ordering, one discovers that norms are modelled everywhere outside the field of Artificial Intelligence and Law: In

⁵Nulla poena sine lege, after all

planning and design systems they usually occur as requirements (that must be met) and constraints (that must not be violated).

We have studied pragmatic automated normative reasoning in detail in earlier projects (generally [22, 4, 26]). We treat norms as logical sentences describing a situation and assigning a value of ‘allowed’ or ‘disallowed’ to it. We (cf. [26]) simply replace obligations with the equivalent prohibition (using the standard observation that $O(A \rightarrow B)$ is the same as $F(A \wedge \neg B)$) and the weak permission $P(A \wedge B)$. This allows us to treat obligations as value assignments useful for assessment of situations.

In the context of case-based reasoning (e.g. [1]) it is commonly acknowledged that many deontic norms that distinguish between only two valuations can be considered indicative of the existence of a more fundamental preference of the legislator underlying it. If the legislator for instance orders civil servants to always buy the cheapest of some good, that statement is indicative of a more general preference for cheaper goods over more expensive goods. In court it is often possible to argue successfully that this underlying scale represents the actual intent of the legislator and can be taken into account (see [1] for some wellknown examples).

If this is the case, one wonders why the legislator is not more direct in the expression of his preferences. The best explanation is that continuous scales are complex to handle in case of a conflict between several competing values, but we have already noted that this is what will happen in practice in court rooms. Another possible explanation is the attractiveness of clearcut good-bad distinctions instead of better-worse scales. A third option is that the legislator assumes a general lack of mathematical skills among the users of legislation⁶.

3.1 Norm Systems and Conflict Resolution

Norms establish a preference ordering on situations to determine what one ought to do given a choice between two situations, but they do not intend define a generic recipe for establishing a total preference ordering on situations that for instance requires one to plan a legally perfect life for oneself, or that allows the DTCA to forecast the exact effect of legislation on taxpayer’s choices. The legislator limits itself to surgically changing existing preference relations with norms that assign a value to an incomplete description of a situation, assuming that the rest of the situation is irrelevant to its value.

Competing valuations of the same situation, or conflicting norms, require a valuation of the conflicting norms, and a choice between them or a weighing procedure that takes the relative value of the involved norms into account. The ordering on norms does not imply an ordering on (ideal) worlds; No decision-theoretic assumptions like preferential independence between situations can be made. Violating two norms is for instance not necessarily less ideal than violating one norm. A legislator for instance issues conflicting norms if it orders its civil servants to always accept the cheapest bid for a construction contract, and to see to it that always the most durable materials are used. Because in both norms an implicit general preference is apparent, this conflict can be solved with a weighing procedure. A reasonable choice respecting the intentions of the legislator will likely violate both norms. These norms conflict, but the

⁶Also evidenced by the consistent use of tables instead of formulae in tax legislation. One wonders why these users are expected to understand extremely complex sentences while nonlinear functions are considered too complex.

competing preference scales are compatible.

In most cases, however, conflicting norms are incompatible and a choice is necessary. A fictional example:

1. *Tanks adjacent to the hull are not used to store fuel oil.*
2. *If the ship is fitted with a double hull, tanks adjacent to the hull can be used to store fuel oil.*

In regulations these exceptions are often used for compactness; A so-called qualification model (a full, monotonic paraphrase, cf. [7]) of regulations is often much harder to understand and use. If explicit and intended, the exception relation is often marked in the text in some way. Usually the more specific regulation is superior, unless it is of lower order or of equal order but issued earlier (both properties of the containing regulation). These principles of choice impose a priority ordering on norms (not situations) and go by the names *Lex Specialis*, *Lex Superior*, and *Lex Posterior* respectively in legal theory.

To ‘solve’ the valuation of a situation description, the situation description must be completely available: all relevant facts must be there. If you move a case for which you already know the qualification from one jurisdiction to another, it cannot be assumed that all necessary facts are supplied at the right level of detail, because you cannot determine in advance what the relevant facts are. The case description is focused on matching a specific legal vocabulary, and which facts are supplied is based on the presuppositions made about what is relevant to the normative qualification of the presented case. It is therefore in principle impossible to rejudge a case in another jurisdiction based on a written verdict.

4. HARMONIZATION

With a theory of how legislation supposedly influences human behaviour, we can explain what harmonization in the EU usually aims at [15]:

Legal Convergence aims for similar legal responses to the same situations.

A ‘similar’ legal response can mean a lot of things: If a fine for a certain traffic violation in the Netherlands is higher than a fine in a similar situation in Portugal, they may still be qualitatively the same. It depends on the characteristics of legal responses that are we interested in. For the civil law jurist it is often enough to establish that both countries give a fine for the same behaviour. For the economist it is the amount of the fine that determines how much it affects behaviour. It depends on the norms of analysis used. Taxes are usually analysed as an ‘interference’ in normal economic behaviour, abstracting away completely from the way they are collected and why they are collected that way. This ‘economic’ view fails to distinguish between the cost of a fine and the cost of taxes, even though it is commonly accepted in legal economics that the reputation cost of violating the law to a person or organization can be empirically established by quantifying the costs a person or organization is willing to make to avoid the violation or hide the violation from others. Reasonable candidates for legal responses include deontic qualifications, public acts, reparations ([15]; in our view rather specific for acts that have an apparent victim), fines, prison sentences, etc.

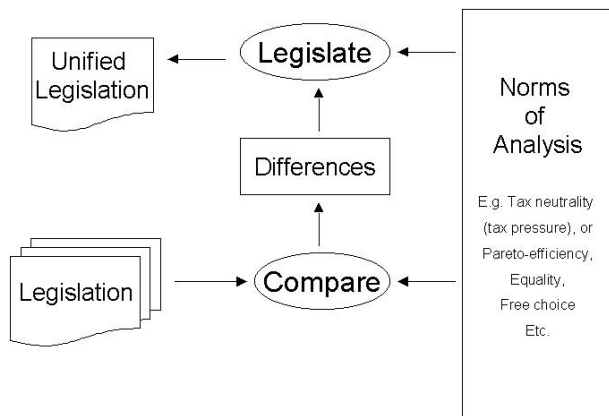


Figure 1: The harmonization process

4.1 Legislative Strategies for Harmonization

There are two general strategies for achieving legal convergence. The first is *unification*, or replacing ‘similar’ but different norms with one norm. The other one is *standardization*, or issuing general norms of analysis prescribing what effects national legislation should achieve. Unification can also be achieved without multilateral agreement through coercive imposition or reception [15]. In the European Union the following classification of legal documents is common:

Regulation Adoption of a regulation (classification CEN EN) makes a regulation directly applicable horizontally and vertically. It impose an obligation of form on everyone.

Directive Adoption of a directive (or harmonization document; classification CEN HD) is binding with respect to common aims, but leaves implementation to the adopting legislators. It imposes an obligation of result on the legislator. In some cases, directives or treaty provisions are considered *self-executing* in some jurisdictions because of their form; Rights can be (vertically at least – horizontal is matter of debate) implied by obligations imposed on legislators if the directives are clear and unconditional (wrt. a case).

Decision A decision only binds the parties addressed.

Figure 1 shows a inference structure that relates the input regulations, differences detected, norms of analysis, and unified legislation. We propose that a European regulation (CEN EN) corresponds in Figure 1 with unified legislation, while a European directive, or harmonization document (CEN HD), posits norms of analysis for the box on the right hand side. Adding an input arrow from the input legislation to the legislate task describes the situation where existing legislation is amended to comply with a directive.

4.2 Norms about Norms

Norms intended for legislators also exist; European ‘directives’ impose a duty on legislators to change their legislation to achieve certain aims. Whether a legislator violates a directive usually appeals to the perceived intentions of the legislator, or the perceived causal effect of the norms on the aims to be achieved – which is difficult to establish.

Norms of analysis are the norms used to distinguish good from bad norms (cf. generally [13, 12] in taxation context), and represent the values the legislator is committed to. We have noted that courts also appeal to ‘values’ or ‘intentions’ to judge the propriety of applying a norm ([1]). These values can take the form of a continuous or discrete metric. To distinguish good from bad legislation, one has to deal with the additional problem of aggregation of norms in regulations, and the aggregation of the results of application of multiple norms of analysis. Tax neutrality is such a norm of analysis that postulates that taxation should not create avoidance behaviour. Tax neutrality is also a norm for distinguishing good and bad combinations of legislation in the context of migration; In this form it postulates that two jurisdictions that allow free movement between jurisdictions should try to minimize differences in tax pressure. The respective tax laws of both jurisdictions are ‘tax neutral’ towards eachother if they give no incentives or rewards for tax avoidance by movement between jurisdictions. The alternative is tax competition. Most general political values (like Pareto-efficiency, solidarity, deregulation, subsidiarity, the ‘free market’) can serve as a norm of analysis for distinguishing a good from a bad law.

We do not subscribe to the view that $O(OA) \rightarrow OA$ (the deontic axiom that if it ought to be that it ought to be that B then it ought to be that A) when a norm about a norm $O(OA)$ is a directive from another legislator. As Stephen Brown J. remarked in *Farrall vs. Department of Transport* [1983] RTR 279 (at 291), where Mr. Farrall contended it is allowed to drive a motor car on British roads without a licence as required by section 85 of the Road Traffic Act 1972 :

”It is a misunderstanding that any statute can be regarded as null and void because of the EEC Treaty. What is required is that the member state shall introduce regulations or legislation which shall give effect to decisions which are binding because of the Treaty. That of course has now been done.”

In a dualist jurisdiction, where international law is not explicitly automatically part of national law, the international law must be explicitly incorporated in order to have any legal effect. An international law may be phrased in such a way that it suggests direct effect (e.g. the Treaty of Rome) but it in effect imposes duties on legislators. Monist jurisdictions (like the Netherlands, constitutionalized in art. 93, 94 Gw and HR 3 maart 1919, NJ 1919, 371, Grenstraktaat Aken) that directly accept provisions with direct effect are rare. In a monist jurisdiction it is more or less true that if it ought to be that it ought to be that A then it ought to be that A for provisions that can have direct effect, but the provision that allows for this is explicitly stated in legal sources. In such a case, directives or treaty provisions are considered self-executing because of their form.

We have distinguished three categories of norms about norms with different strength:

Norm of analysis A continuous or discrete metric that imposes an ordering on descriptions of norm systems that represents the degree of desirability of the world covered by the description.

Directive A norm of analysis to which a legislator has committed itself. The legislator legislates in accordance with the directive and accepts the ordering on ideal worlds implied by it.

Self-executing Directive A two-value deontic directive for which the axiom $O(OA) \rightarrow OA$ is accepted. Direct effect is lim-

ited to deontic norms, because for more complex metrics it is not clear how it should translate to behaviour.

More interesting than the relatively simple ‘deontic’ norms about norms are more general norms that refer to different scales to order worlds in degrees of ideality. To operationalize these norms one often has to make sweeping generalizations of the effect of deontic norms on behaviour. In the harmonization context these norms are often explicit in the directives, in other contexts of comparison they exist but usually remain implicit. The ‘commonsense’ perspective on income tax systems, for instance, treats it as a black box where your income in the input and the part of your money you may keep for yourselves is the output. What happens here is that the analyst who compares these black boxes makes a normality assumption about behaviour: all civil servants do what they ought to do and you have honestly submitted all relevant information. $O(A \rightarrow B)$ no longer translates to $F(A \wedge \neg B)$ but to $A \rightarrow B$ for all norms irrelevant to the analysis. The analyst relying on this ‘commonsense’ perspective behaves like the sociologist in section 3 and *describes* the actions of a tax administration as if the relevant tax legislation is a specification of its behaviour (cf. generally [2]). A simple comparison of black boxes is then sufficient to rank tax systems according to for instance marginal tax pressure.

5. REPRESENTING NORM SYSTEMS AS ONTOLOGIES

To make the integration of norm systems fit the paradigm of ontology integration [6] we have to treat the norm system as an ontology, and we have to restrict our interface to our norm system so that it is the same as that of a description classifier. In the past (cf. [23, 22]) we tested the violation of norms by querying the interface of the description classifier for instances of the situation description described in the norm. If we represent norms as concepts it is clear that the valuation attached to the norm is not monotonic – because of the exceptions – and a higher order property that applies to the norm as object, not a property of the ‘instances’ of the norm. Conceptually we can say that there are two subtypes of the ‘concept’ concept: the prohibition and the permission. Each prohibited or permitted generic case is an instance of these concepts. Valuation scales with more than two values can be treated in the same way.

UML, the modeling language used in the E-POWER project, is – despite its aspirations to being an international standard for models – not suitable for ontology integration. It lacks a description classifier and suffers from some other limitations that make it unsuitable for expressing norms in such a way that they can be classified. Instead we use the W3C’s Web Ontology Language (OWL)⁷, the successor to DAML+OIL and the proposed representation language for the W3C’s ‘Semantic Web’. OWL is the standard ontology extension to be used with RDF⁸ and can be embedded in METALex XML documents.

5.1 Norm systems

For a more precise notion of what we consider a norm system, let \mathcal{R} be the set of modelled regulations, \mathcal{C} the resulting set of cases, set $\mathcal{Q} = \{allowed, disallowed, silent\}$, and the relation $\mathcal{N} \subset \mathcal{R} \times \mathcal{C} \times \mathcal{Q}$ the set of norms. Cases are expressed as closed first-order formulas, built inductively using monadic or dyadic predicates P_1, P_2, \dots, P_n , variable letters x, y, z, \dots , and logical con-

nectives \neg and \wedge . All variables are existentially quantified and distinct variable letters are assumed disjoint within one formula. A predicate letter denotes a term from the domain ontology. Normalized to a cube (conjunction) of literals a case $\varphi \in \mathcal{C}$ is written as:

$$\varphi = [P_1x_1, P_2x_2, \dots, P_nx_n]$$

A case φ entails another case ς iff any interpretation using the domain ontology that makes φ true also makes ς true. If the user supplies a case ς for valuation, there is a norm $(r, \varphi, \rho) \in \mathcal{N}$, and $\varsigma \models \varphi$, then obviously $(r, \varsigma, \rho) \in \mathcal{N}$. If, for example user case $\varphi = [ship(x), doubleHull(y), fuelOilTank(z), fittedWith(x, y), fittedWith(x, z), adjacent(y, z)]$ and the previous example regulations in section 3.1 are identified as r_1 and r_2 then it is possible to infer $(r_1, \varphi, disallowed)$ and $(r_2, \varphi, allowed)$ if the domain ontology is sufficiently complete. A norm may be an *exception* to another norm if it strictly disaffirms the other norm as defined here:

DEFINITION 1. (Strict Disaffirmation) A norm $(r_1, \varphi, \rho_1) \in \mathcal{N}$ strictly disaffirms a norm $(r_2, \varsigma, \rho_2) \in \mathcal{N}$ iff. $\varphi \models \varsigma$ and $\varsigma \not\models \varphi$ and $\rho_1 \neq \rho_2$.

This relation between norms can be computed offline and stored in a noncircular partial ordering $\mathcal{E} = \{\mathcal{N}_2, \prec\}$. The example norms are thus resolved to $(r_1, \varphi, disallowed) \prec (r_2, \varphi, allowed)$ and φ is allowed if r_1 is not of higher order or issued later and of equal order. The resulting exception graph is input to a bigger noncircular partial ordering $\mathcal{P} = \{\mathcal{N}_2, \prec\}$ which is the overall priority ordering on norms used to establish whether a case is allowed or disallowed by \mathcal{P} , or whether \mathcal{P} is silent about the case. The other ordering principles do not depend on subsumption relations between situations.

We have in the past described a simple strategy for expert systems of continually trying to apply norms that change the verdict from allowed to disallowed, or vice versa (cf. generally [26, 3]) to obtain all relevant information from a lay user. The order in which norms are tried is based on the exception graph. The presentation of potential exceptions can be fine-tuned further by determining the prime implicant of the conjunction of all missing facts (to remove logical redundancy in the line of questioning; cf. [21, 17]) and making smart ‘frame’ assumptions (like the assumption that others than yourself always behave in accordance with the relevant regulations).

5.2 Representing Norms in OWL

OWL is an extension to *RDFS* for ontologies that imposes a much stronger interpretation on RDF graphs – in OWL one can define fairly precisely what graph structures are valid. OWL is based on an *ALC* Tbox (terminology box), a set of assertions about concepts of the form $C_1 \sqsubseteq C_2$ (and $C_1 \doteq C_2$ as shorthand for equivalence) where C_1 and C_2 are concept definitions conforming to the syntax $C ::= N \mid \top \mid \perp \mid \neg N \mid C_1 \sqcap C_2 \mid C_1 \sqcup C_2 \mid \forall R.C \mid \exists R.C$ where N is a concept name with the usual interpretation. Roles can be defined with primitive role definitions $N \sqsubseteq R$ where R is a role conforming to the syntax $N \mid R_1 \sqcap R_2$ and N is a role name with the usual interpretation (including frame identifiers for transitivity, inversions, and functions).

To map norm cases to OWL we have to solve the problem of the

⁷<http://www.w3.org/2001/SW>

⁸The Resource Description Framework; <http://www.w3.org/RDF/>

missing existential variables for queries. There is a folding procedure for existential variables in conjunctive queries (cf. [14]). A lot of the reasoning effort can be moved to the OWL description classifier and we know that the only errors remaining will be false positives. The general idea is fairly simple. Let $Case \sqsubseteq (Ship \sqcap \exists FW : ME_1 \sqcap \exists FW : (FP \sqcap \exists DB : ME_1) \sqcap \exists FW : FP)$ be a case of a ship fitted with a main engine and two fire pumps, one driven by that main engine, $ME_1 \sqsubseteq MEng$ an instance of a main engine, and $N_1 \doteq (Ship \sqcap \exists FW : ME \sqcap \exists FW : (FP \sqcap \exists DB : MEng))$ and $N_2 \doteq (Ship \sqcap \exists FW : FP_1 \sqcap \exists FW : FP_2)$ be norms, where $FP_1 \doteq FP$ and $FP_2 \doteq FP$ are variables that can be matched with an instance of a fire pump. As expected, $Case \sqsubseteq (N_1 \sqcap N_2)$.

The problem is that $(N_1 \sqsubseteq N_2)$ is also true (because $FP_1 \doteq FP_2 \doteq FP$): We can enforce co-reference, but we cannot force FP_1 and FP_2 to match different instances of firepumps. Adding to the Tbox that $(FP_1 \doteq (\neg FP_2))$ merely makes the Tbox inconsistent. The problem can be solved with an algorithm that obtains the set of generic cases that subsume the case, and for each generic case obtains the facts subsumed by each variable in the generic case, and finds a mapping from all variables to distinct instances in the case. A slight disadvantage is that it requires all instances (even in different cases) to have a *unique* identifier, but that is also a RDF requirement for resources. This is why RDF toolkits usually include a unique URN generator. The exception graph defined in the previous section and [26, 3] now corresponds with the subsumption graph in the description classifier, except for the category of the norm – prohibition or permission.

6. COMPARING ONTOLOGIES OF LEGISLATION

To compare legislation, we initially have the choice of what legislation to compare. Since it is unfeasible to compare complete legal systems, we have to make an educated guess about whether regulations address similar situations, and whether they are the only regulations that address those situations. For the Dutch and Belgian Income Tax Law that is fairly obvious; Mismatches mostly concern situations regulated in other laws that must be involved. Comparing the Dutch General Administrative Law with a British equivalent is impossible; Since the latter does not exist it is better in this case to start from a ‘space’ of possible situations, and find pieces of legislation and court verdicts that address similar situations.

Only situations covered by both models can be compared. The differences we are interested in, are inferred from ‘similar’ situations that lead to qualitatively different outcome valuations (e.g. different amount fine, marginal tax pressure) in the norm system we are comparing with.

6.1 Explicitly assigning Similarity Relationships

The ontologies of concepts in legislation currently used by the DTCA very closely follow the text of tax legislation – an approach that improves intercoder reliability for modelers. Liberal application of commonsense additions the model is taboo because it is too subjective. For the design of decision support systems this need not be a practical problem, because software designers usually assume that ‘extensional’ commonsense connections, for instance the subsumption relation between senior and adult (all seniors are adults), will be transparent to the user of the system. But if you want to

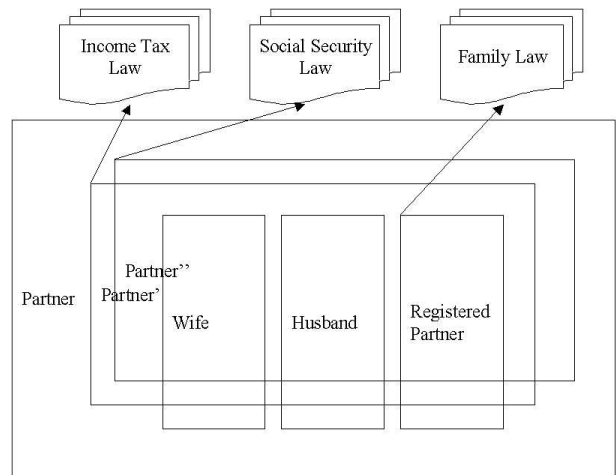


Figure 2: Legislation and models of Legislation

compare two models of different income tax laws, you must find a semantic grounding in shared ‘commonsense’ concepts to make them comparable.

In the current POWER method, modelers are required to model each fragment of law, and the partial definitions of concepts occurring there, separately. In the integration phase, the modeler may decide that two occurrences of the same phrase in different fragments of law, for instance ‘partner’, refer to the same concept. One of the purposes of making these choices explicit is to be able to communicate the differences between different meanings of a concept like ‘partner’ to the end user of decision support systems. Comparing two different regulations from different jurisdictions requires another integration phase. The integration decisions can be added as separate statements of similarity, including an annotation who made the decision (the modeler, DTCA policy, the fiscal judge etc.).

Figure 2 shows the number of different extensional assumptions that must be made to establish for instance a ‘similarity’ between different some interpretations of ‘partner’ in the family context to make them comparable. The Venn diagram displays three definitions of ‘partner’ from (Dutch) law, and the common sense concept as it is interpreted. Since the assumptions can all be wrong, it is the responsibility of the modeler to judge whether differences are real or artifacts caused by interpretation. A new concept ‘partner’ can cover or overlap any of the 13 extensional spaces currently in Figure 2.

6.2 Ceteris Paribus Assumptions

Because we can only compare a limited space of possible situations, we have to make the assumption that the rest of the situation is for purposes of the comparison irrelevant. The fact that a situation occurs in Amsterdam or London, or today or yesterday, or that it involves other persons is obviously not a relevant difference for the purposes of comparing legislation. But suppose that we want to compare the Dutch and British traffic code, which both regulate the ‘same’ traffic situations. We model the correspondences between Dutch and British to find out that the only thing that a formal reasoning procedure can prove to us is that every situation allowed in the Netherlands is disallowed in the United Kingdom, and there is no other meaningful comparison to be made. A useful and reasonable frame assumption to make in these circumstances is that ‘right

in Netherlands is equivalent to left in England” for the purposes of the analysis, because driving right is normal in the Netherlands, and driving left is normal in the UK. This set of assumptions added to seal of the situation space are the *Ceteris Paribus* (other things being equal) assumptions:

Ceteris Paribus Assume that the rest of the situation you are comparing is equally ‘normal’ in both jurisdictions. The situation that is compared is normal in both jurisdictions if the participants in the situation value the situation similarly (in terms of utility, morality, likelihood etc.).

In some cases we translate situations to accommodate general patterns of preference holding in a specific jurisdiction, to make them equivalent in valuation. This is intuitively right; Everyone agrees that although driving left may constitute a crime in a certain jurisdiction, the behaviour itself is morally neutral (as opposed to murder, for instance). The important thing is that these assumptions are made explicit in a separate assumption model, that contains our assumptions about the ‘transfer’ of concepts and value judgements. This model contains any definition that does not occur as such in the legislation that is compared, but is added by the modeller to explain the perceived links between legislation. The normality assumptions about behaviour (most common: all others do what they ought to do) belong in this class.

6.3 Norms of Analysis

We also have to define the preference ordering or weighing procedures for norms of analysis that can also be seen as *ceteris paribus* comparatives (cf. generally [8]). Since these are norms about norms they can only be directly applied to classified cases and tell us whether the classifications are good or bad if the necessary normality assumptions have been defined. A norm is after all a metric that assigns a value to the incomplete description of the world that the situation description represents assuming that the rest of the world that is not described is morally neutral. If assigned values ‘compete’ for the same case, we also need a conflict resolution method that chooses between norms of analysis, or weighs them.

6.4 Explaining Relevant Differences

When relevant differences in valuation are found, the modeler has to discover what caused the differences. It may be the legislation that is qualitatively different, but it may also point to a different interpretation of the modelers of both legislations, or a suspect assumption in the comparison. An alternative problem may be that we are comparing the wrong set of legislation. Fortunately, all of these findings have the potential of either improving the legislation, or its ontology. A phone call or email message to experts involved in applying the regulations is often sufficient to attribute the result.

7. DISCUSSION

We have argued in this paper that any comparison of norms (implicitly) involves norms of analysis. To assign a value to a norm, one has to make assumptions about the effects of the norm. We have also argued that comparison of limited slices of a legal system requires assumptions about what is the same, and the assumption that everything left out of consideration is irrelevant. In a ‘Semantic Web’ context we can think of these packages of assumptions, annotated with the identity of the authority that made them available, as reasoning modules that allow us to draw interesting conclusions if we trust the authority.

We understand from small experiments that the amount of external information from the legal system needed to explain in what way two regulations are different, is potentially very large. Experimenting with comparative analysis on a somewhat larger and methodical scale will teach us whether these extra requirements increase explosively or level off at some point due to discovery of regular patterns. We are currently comparing some very similar pieces of Dutch legislation (tax and social security), and will try some Dutch and Belgian legislation (pensions) soon. We do note that making simplistic assumptions can always help to keep things manageable.

An implicit assumption underlying the notion of comparing knowledge models of legislation is that the knowledge models are complete declarative statements of the ‘normative’ reading of legislation, that have been integrated with, or translated to business processes afterwards. In large organizations like the DTCA, the relations between legislation, business processes, and IT infrastructure are sufficiently complicated to appreciate the advantages of principled normative models of legislation for designing and maintaining IT infrastructure and business processes. Being able to compare legislation is an added value of models that are valuable per se.

To make knowledge models of legislation with the purpose of comparing legislation from different jurisdictions is a huge additional effort that is often not necessary to implement practical decision support systems. At the same time, if you fail to make the effort to ground your knowledge model in commonly understood concepts (concepts for which we can point out the ‘extension’, or the things to which it applies) it becomes as difficult to discover the interactions between different regulations in the same jurisdiction, or even one tax administration, as it is to compare regulations from different jurisdictions. Comparing with foreign law is useful for that reason alone: It is generally not very convincing to show that the computer can make ‘smart’ commonsense connections between concepts in legislation that is very familiar to its users. In addition, the problem of understanding foreign tax legislation is becoming more and more relevant as more people move over borders and expect for instance to be able to find out whether this will have fiscal consequences for their pensions.

We have also remarked that in translation from legal jargon to ‘commonsense’ vocabulary one must strike a balance between using the concepts an audience already knows, and teaching the definitions of unfamiliar legal concepts. This obviously causes ‘conceptual drift’ (cf. [19]) of commonsense concepts to include the legal concept of the same name. This is happening across legal cultures in the EU. We have not explored how this factor may affect harmonization and the stability of knowledge models.

Subjects like harmonization of taxation are very complex, and for that reason often presented as purely ‘technical’. Questionable political choices in such areas can remain undetected for a long time. That is why merely exposing the values that underpin EU guidelines is an active field of scientific inquiry worthy of publication (see for taxation and pensions for instance [9]). The Knowledge Engineering approach to comparing legislation presented in this paper is a more systematic and transparent approach to analysing and documenting the development of legislation than the inherently political ‘functional’ approach a human analyst has to resort to. It separates the value systems applied from ‘similarity’ relationships and has the potential to improve the documentation and motivation of value judgements embedded in it.

8. ACKNOWLEDGEMENTS

E-POWER is partially funded by the EC as IST Project 2000-28125; partners are the Dutch Tax and Customs Administration, O&I Management Partners, LibRT, the University of Amsterdam (NL); Application Engineers, Fortis Bank Insurance (B); Mega International (F).

9. REFERENCES

- [1] T. Bench-Capon and G. Sartor. Using values and theories to resolve disagreement in law. In J. Breuker, R. Leenes, and R. Winkels, editors, *Legal Knowledge and Information Systems (JURIX-2000)*, Amsterdam, 2000. IOS Press. ISBN 1.58603.144.9.
- [2] T. J. M. Bench-Capon. Deep models, normative reasoning and legal expert systems. In *Proceedings of the second international conference on Artificial intelligence and law*, pages 37–45. ACM Press, 1989.
- [3] A. Boer. The Consultancy Game. In J. Breuker, R. Leenes, and R. Winkels, editors, *Legal Knowledge and Information Systems (JURIX-2000)*, pages 99–112, Amsterdam, 2000. IOS Press.
- [4] A. Boer. MILE: Turning Legal Information into Legal Advice. In A. Tjoa and R. Wagner, editors, *Proceedings of the Twelfth International Workshop on Database and Expert Systems Applications (DEXA)*, pages 787–791, Los Alamitos (CA), 2001. IEEE Computer Society.
- [5] A. Boer, R. Hoekstra, and R. Winkels. METALex: Legislation in XML. In A. D. T. Bench-Capon and R. Winkels, editors, *Legal Knowledge and Information Systems. JURIX 2002: The Fifteenth Annual Conference.*, pages 1–10, Amsterdam, 2002. IOS Press.
- [6] D. Calvanese, G. D. Giacomo, and M. Lenzerini. A framework for ontology integration. In *Proceedings of the First Semantic Web Working Symposium*, pages 303–316, 2001.
- [7] N. den Haan and J. Breuker. Constructing Normative Rules. In *Proceedings of JURIX'96*, pages 135–147, 1996.
- [8] J. Doyle and M. P. Wellman. Representing preferences as ceteris paribus comparatives. In *Decision-Theoretic Planning: Papers from the 1994 Spring AAI Symposium*, pages 69–75. AAAI Press, Menlo Park, California, 1994.
- [9] E. Engelen. Financialization, Pension Restructuring, and the Logic of Funding. In *Proceedings of SASE 2002*, 2002.
- [10] A. Gangemi, N. Guarino, C. Masolo, A. Oltramari, and L. Schneider. Sweetening Ontologies with DOLCE. In A. Gómez-Pérez and V. R. Benjamins, editors, *Knowledge Engineering and Knowledge Management. Ontologies and the Semantic Web, 13th International Conference, EKAW 2002, Sigüenza, Spain, October 1-4, 2002*, volume 2473 of *Lecture Notes in Computer Science*, pages 166–181. Springer, 2002.
- [11] A. Gangemi, D. M. Pisanelli, and G. Steve. A formal Ontology Framework to Represent Norm Dynamics. In *Proceedings of the Second International Workshop on Legal Ontologies (LEGONT)*, 2001.
- [12] W. Hettich and S. L. Winer. Economic and political foundations of tax structure. *American Economic Review*, 78(4):701–712, 1988.
- [13] W. Hettich and S. L. Winer. Rules, Politics, and the Normative Analysis of Taxation. Carleton Economic Papers 00-12, Carleton University, 2000.
- [14] I. Horrocks and S. Tessaris. A conjunctive query language for description logic aboxes. In *AAAI/IAAI*, pages 399–404, 2000.
- [15] M. Kohen. Europe and the Standardization of the Law: Past and Present. In *Proceedings of the Conference on Globalisation and International Relations in the 21st Century*, 2002.
- [16] M. Muller. The RDF Dictionary and the standardisation process in the legal domain. In *Proceedings of XML Europe*, 2002.
- [17] L. Pallipoli, F. Pirri, and C. Pizzuti. Algorithms for Selective Enumeration of Prime Implicants. *Artificial Intelligence*, 111:41–72, 1999.
- [18] H. Pinto, A. Prez, and J. Martins. Some issues on ontology integration. In *Proceedings of the IJCAI-99 Workshop on Ontologies and Problem-Solving Methods (KRR5)*, 1999.
- [19] E. Rissland and T. Friedman. Detecting change in legal concepts. In *Proceedings of the Fifth International Conference on Artificial Intelligence and Law (ICAIL-99)*, pages 127–136, New York (NY), 1995. ACM.
- [20] S. Spreeuwenberg, T. van Engers, and R. Gerrits. The Role of Verification in Improving the Quality of Legal Decisionmaking. In B. Verheij, A. Lodder, R. Loui, and A. Muntjewerff, editors, *Legal Knowledge and Information Systems (JURIX-2001)*, pages 1–16, Amsterdam, 2001. IOS Press. ISBN 1.58603.201.1.
- [21] J. Straach and K. Truemper. Learning to Ask Relevant Questions. *Artificial Intelligence*, 111:301–328, 1999.
- [22] A. Valente. *Legal Knowledge Engineering: A Modeling Approach*. PhD thesis, Amsterdam, 1995.
- [23] A. Valente and J. Breuker. ON-LINE: An architecture for modelling legal information. In *International Conference on Artificial Intelligence and Law (ICAIL-1995)*, pages 307–315, 1995.
- [24] L. W. N. van der Torre and Y.-H. Tan. Contextual deontic logic. In P. Bonzon, M. Cavalcanti, and R. Nossum, editors, *Formal Aspects of Context*, pages 143–160. Kluwer Academic Publishers, Dordrecht, 2000.
- [25] C. van Laer. The Applicability of Comparative Concepts. *European Journal of Comparative Law*, 2.2, 1998.
- [26] R. Winkels, D. Bosscher, A. Boer, and J. Breuker. Generating Exception Structures for Legal Information Serving. In T. Gordon, editor, *Proceedings of the Seventh International Conference on Artificial Intelligence and Law (ICAIL-99)*, pages 182–195, New York (NY), 1999. ACM.