

A Knowledge Engineering Approach to Comparing Legislation

Alexander Boer¹ and Tom van Engers²

¹ Dept. of Computer Science & Law, University of Amsterdam,

² Dutch Tax and Customs Administration, Utrecht, Netherlands

Abstract. In the E-POWER project relevant tax legislation and business processes are modelled 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. Better insight in the process of modeling and comparing legislation from different legislators is expected 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, 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. This paper discusses known problems and requirements of comparing legislation, and the expected results of comparing models of legislation.

1 Introduction

In the E-POWER project relevant tax legislation and business processes are modelled 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 conceptual models 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 [16].

We are currently researching whether conceptual models can also be used to compare ‘similar’ legislation from different jurisdictions. 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.

Comparing similar regulations from multiple jurisdictions is 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 notable 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 usually have very different meanings and are written in completely different styles. Different jurisdictions, cultures, and languages, and the philosophical problems of comparing legal systems are the main subject of Comparative Law. Researchers in Comparative Law rarely compare written legislation directly.

But similar regulations are habitually compared in a number of contexts, and not all of them are perceived as equally complex and subjective. Regulations are compared for very different purposes, for instance:

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 combines these purposes; 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.

There are a number of initiatives for constructing international legal ontologies that expose the subsumption relations between legal vocabulary in multiple jurisdictions. Gangemi et al. [9] suggest, in a context of comparing versions of a regulation, that the problem of comparing legislation is a special case of the more general problem of ontology integration (cf. [4]). This paper discusses some the problems we run in to when we try to translate the problems of Comparative Law into those of ontology integration.

The next sections discuss how to establish ‘similarity’ between concepts, some concepts needed to describe and explain legislation itself, and the specific problems related to comparing knowledge models of limited scope. The expected results of comparing knowledge models of legislation will be discussed in the last section.

2 Similarity of Concepts in Legislation

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 Dutch law and society. This problem is not unique for translation of *legal* texts, but it is especially obvious and 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.

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 competence. 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 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 purpose of the translation.

Intuitive ‘similarity’ can for purposes of Comparative Law be classified as extensional, immanent, or functional [19]. 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* 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 this sense concepts can be functionally ‘similar’, even though the definitions are different.

Because the similarity is based on either the perceived causal effect of the concept on the ‘similar problem’, or the perceived intentions of the legislator (the intended causal effect), 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 implicit value systems and assumptions about causal 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, these implicit underlying values that make things functionally ‘similar’ must be made explicit.

3 Concepts describing the Role of Legislation

To explain why one regulation is better than the other, and how regulations affect the real world, we need concepts that explain the role that the text of the regulation plays in a jurisdiction. One example of generic legal concepts for the design of legal expert systems is the *Functional Ontology of Law* of Valente (cf. [17]), also one of the first to stress the importance of semantically grounding legal concepts in commonsense concepts.

The key concept for understanding legislation is the *norm* (e.g. [6, 17]); 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 are usually made explicit in a document (contract, regulation, jurisprudence, etc). The document often prescribes behaviour to agents assigned a certain role, for instance the owner and user of a road. Sometimes the document posits norms for the creation of artefacts – things made by agents, for example ships or tax forms. In other cases it constrains or defines procedures for actions or transactions by agents, for example a survey, hearing, or purchase, instead.

In deontic reasoning systems norms are treated as logical sentences describing situations that are *allowed* or *disallowed*. They mainly serve as an aid in establishing a preference ordering on situation descriptions to determine what one ought to do given a choice between situations. Because norms in legislation define a discrete metric with just two values (allowed, disallowed) there is no generic recipe for establishing a *total*

preference ordering on situations that decides every choice one is confronted with in life, or that allows the DTCA to forecast the exact effect of legislation on taxpayer's choices for instance. Common decision-theoretic assumptions like preferential independence between disallowed situations do not necessarily hold for legislation. It is therefore in principle also impossible to predict the effects of legislation on choices. We have studied pragmatic automated normative reasoning in detail in earlier projects (generally [17, 2, 20]).

Deontic logics are logics of ideal worlds. Ideal worlds tell us little about what to do. If deontic sentences are interpreted as logical sentences describing an ideal world, they cannot be used to express what ought to be done in a sub-ideal world; This is the essence of the so-called Chisholm and good Samaritan paradoxes (e.g. [17]). It shows that deontic sentences can presuppose a sub-ideal world. Norms in legislation often express competing underlying values that may prove to be jointly unrealizable in some situations. In addition, some norms are (intended as) exceptions to other norms. In regulations these exceptions are often used for compactness; A so-called qualification model (a full, monotonic paraphrase, cf. [6]) 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. In principle the more specific regulation is superior, unless it is of lower order or of equal order but issued earlier. These principles impose a priority ordering on norms (not worlds) and go by the names *Lex Specialis*, *Lex Superior*, and *Lex Posterior* respectively in legal theory.

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. The application of codified norms by a judge to others is also regulated by norms. Norms intended for legislators also exist; European 'directives', for instance, impose a duty on legislators to change their legislation to achieve certain aims. Whether a legislator violates a directive again appeals to the perceived intentions of the legislator, or the perceived causal effect of the norms on the aims to be achieved – which is often impossible to establish.

We use the phrase 'norms of analysis' for norms used to distinguish 'good' from 'bad' norms (cf. generally [11, 10] in taxation context) that represent the values the legislator is committed to. Courts also appeal to 'values' or 'intentions' of the legislator to judge the propriety of applying a norm (generally [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 each other 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, liberalization, subsidiarity) can serve as a norm of analysis for distinguishing

a good from a bad law. To operationalize norms of analysis, they have to assign a value to situation descriptions.

3.1 The Meaning of Harmonization

Having explained how we think legislation is supposed to influence human behaviour, we can explain what harmonization in the EU usually aims at:

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 and the metric we apply. For the jurist it is often enough to establish that both countries judge a situation as disallowed and give a fine to the actor held accountable for it. For the economist it is the amount of the fine that determines how it affects behaviour. Taxes are usually analysed as an 'interference' in economic behaviour, abstracting away completely from the way they are collected and why they are collected that way. The result of the analysis is operationalized as a quantitative value ordering tax systems according to, for instance, marginal or average tax pressure.

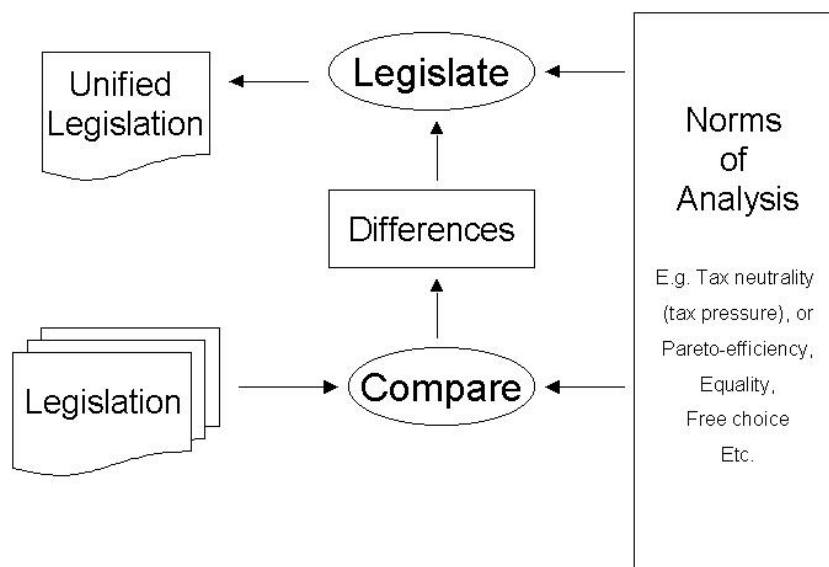


Fig. 1. The harmonization process

There are two general strategies for achieving legal convergence. The first is *unification*, or replacing 'similar' but different norms with one norm. In the EU this is achieved with a regulation, or CEN EN classified document. The other one is *standardization*, or

issuing general norms prescribing what effects national legislation should achieve. In the EU this is achieved with a directive, or CEN HD classified document.

What should be standardized, and to what extent, is a question of which norms of analysis to apply. Figure 1 shows the relations between the differences detected, norms of analysis, and changes to legislation. Note that the norms of analysis play a role twice. Once in determining the differences that are 'relevant' and subject to unification, and a second time in determining what legislation is introduced as a replacement. A EU regulation is an example of unified legislation, and a EU directive posits norms of analysis legislators commit themselves to.

4 Models of Legislation

The UML models currently used by the Dutch Tax and Customs Administration very closely follow the text of tax legislation because that improves intercoder reliability for modelers. Liberal application of commonsense additions to 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 compare two models of different income tax laws, you must find a semantic grounding in shared 'commonsense' concepts to make them comparable.

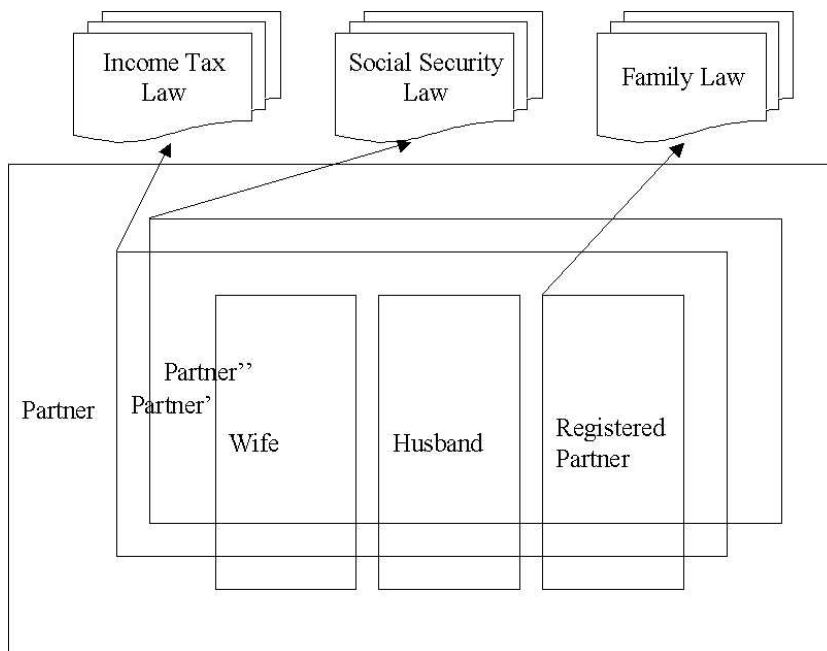


Fig. 2. Legislation and models of Legislation

For the computer, establishing that situations in two different models are similar or different requires a description of the situations using the same concepts, or a *transformation* model containing assumptions that explain the ‘common ground’ between concepts from two regulations. A relevant ‘difference’ may thus be a modeling error in the ‘common ground model’ assumptions. Figure 2 shows the number of different extensional assumptions that must be made to establish a ‘similarity’ relationship between different some interpretations of ‘partner’ to make them comparable. The Venn diagram displays three definitions of ‘partner’ from law, and the common sense concept. 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. Assigning the blame for a difference between norms is not a trivial task in the ontology integration perspective of [9].

In the current E-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. This procedure is needed for integration of concept definitions in one field of law in one jurisdiction. To make models of regulations of different fields of law, and even multiple jurisdictions, comparable the same procedure is required. In this case we need more guidance for classifying concepts by a top ontology because we lose the guidance Dutch administrative law and guidelines for legislative drafting offers for classifying agents, actions, powers, procedures, and principles. To assign the blame for a difference it is also essential to know whose opinion is represented by an assertion in a knowledge base: a taxpayer, knowledge engineer, fiscal expert, the DTCA, fiscal court, or the Court of Justice of the EU?

UML is used because it is widely known among software engineers and supported by familiar CASE tools. UML can be used as a knowledge representation language [5], although it usually lacks a description classifier for verification. We have defined a translation of the exception graph we developed in the past [2, 20] to guide user dialog and knowledge acquisition and to resolve exceptions between norm descriptions to the description classifier paradigm (to be published elsewhere) for OWL (the draft Web Ontology Language; <http://www.w3.org/2001/SW/>) using the ‘folding’ procedure described in [12], so that unresolved differences between norms classify as truth-functional contradictions. The problem of comparing models is translated in this way into a problem of ontology integration (cf. [4, 14])

4.1 Comparing Models of Legislation

To compare legislation, we have to add four categories of knowledge about the comparison we intend to make. Initially, we have to make 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 may seem 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 however impossible; Since the latter does not exist it is better in this case to depart from a ‘space’ of interesting possible situations, and find pieces of legislation and court verdicts that address similar situations.

First, only situation descriptions 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 value system we are comparing with. We have to describe the space of situations that we compare and that the conclusions will pertain to.

Second, because we can only compare a limited space of possible situations, we have to make the assumption that the rest of the situation is static. 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 are the *Ceteris Paribus* (other things being equal) assumptions that have to be added to the ontology to make a specific comparison for a specific purpose. The assumptions state that the situation is valued similarly (in terms of utility, morality, etc.), and the situation is similarly likely (in that place, population, 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). These assumptions are added by the modeller to explain the links between legislation and the perspective we choose – in the abstractions we make – when comparing something.

Third, after establishing our space of situations, we have to choose and operationalize our norms of analysis as *metrics classifying the relevant situation descriptions*. This includes the preference ordering or weighing procedures for multiple norms of analysis that can themselves be seen as *ceteris paribus* comparatives; [7]. Obviously, because norms classify limited situation descriptions (ignoring all facts not asked for) they can conflict wrt. a specific situation.

Fourth, applying norms of analysis on a norm system requires *assumptions about behaviour*. 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.

In deontic terms, this means that an obligation $O(A \rightarrow B)$ no longer translates to $F(A \wedge \neg B)$ but to $A \rightarrow B$ for all norms irrelevant to the analysis. The norms become descriptions of behaviour instead of prescriptions. We have already remarked that the decision as to whether a legislator violates a directive (which is a norm of analysis) appeals to the perceived intentions of the legislator, or the perceived causal effect of the norms on the aims to be achieved. The assumptions about behaviour that we make capture these judgments.

The distinction between intended effect on behaviour and actual effect on behaviour 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 as a bill for health problems (if 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, but the economic analyst will have abstracted away that detail at an early stage.

5 Discussion

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. The European LEXML consortium³ has initiated the LEXML RDF Dictionary project, and the OASIS LegalXML working group⁴ has a committee for a very similar dictionary project. The METALex consortium [3]⁵ has a similar ontology subproject. These projects depart from the premise that a taxonomy of legal concepts is the obvious way to explain the relations between concepts from multiple jurisdictions. In [9] 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. [9] do not account for norm conflict resolution strategies for establishing the validity of norms. We want to add such a component. In addition, we claim that the ontology integration view does not account for the purpose of the comparison. To account for the 'functional' similarities that human analysts see between concepts that the description classifier does not, we have to make the underlying assumptions and value systems explicit.

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 poten-

³ <http://www.lexml.de>

⁴ <http://www.legalxml.com>

⁵ <http://www.metalex.nl>; The authors contribute to this consortium.

tially 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.

When relevant differences 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 transformation model. An alternative problem may be that we are comparing the wrong set of legislation. Fortunately, all of these findings have the potential of improving legislation, or the models of it.

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. Construction of a sizable transformation model is expected to result in an inventory of norms of analysis used for taxation legislation analysis and generic analysis of legislative quality and a core ontology of abstract legal concepts that are necessary that can also be used to evaluate the quality of theoretical legal core ontologies (such as [17]).

To make knowledge models of legislation with the purpose of comparing legislation from different jurisdictions is an additional effort that is not necessary to implement 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 not convincing to show that the computer can make 'smart' commonsense connections between concepts in legislation that is very familiar to its users. 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 'functional' approach a human analyst has to resort to. It has the potential to depoliticize the technical aspect of harmonization by separating it from the value systems applied.

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).

References

1. T. Bench-Capon and G. Sartor. Using values and theories to resolve disagreement in law. In J.A. Breuker, R. Leenes, and R.G.F. Winkels, editors, *Legal Knowledge and Information*

- Systems (JURIX-2000)*, Amsterdam, 2000. IOS Press. ISBN 1.58603.144.9.
2. A. Boer. MILE: Turning Legal Information into Legal Advice. In A.M. Tjoa and R.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.
 3. Alexander Boer, Rinke Hoekstra, and Radboud Winkels. METALex: Legislation in XML. In A. Daskalopulu 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.
 4. D. Calvanese, G. De Giacomo, and M. Lenzerini. A framework for ontology integration, 2001.
 5. S. Cranefield and M. Purvis. Uml as an ontology modelling language, 1999.
 6. N. den Haan and J. Breuker. Constructing Normative Rules. In *Proceedings of JURIX'96*, pages 135–147, 1996.
 7. Jon Doyle and Michael P. Wellman. Representing preferences as ceteris paribus comparatives. In *Decision-Theoretic Planning: Papers from the 1994 Spring AAAI Symposium*, pages 69–75. AAAI Press, Menlo Park, California, 1994.
 8. Ewald Engelen. Financialization, Pension Restructuring, and the Logic of Funding. In *Proceedings of SASE 2002*, 2002.
 9. 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.
 10. Walter Hettich and Stanley L. Winer. Economic and political foundations of tax structure. *American Economic Review*, 78(4):701–712, 1988.
 11. Walter Hettich and Stanley L. Winer. Rules, Politics, and the Normative Analysis of Taxation. Carleton Economic Papers 00-12, Carleton University, 2000.
 12. Ian Horrocks and Sergio Tessaris. A conjunctive query language for description logic aboxes. In *AAAI/IAAI*, pages 399–404, 2000.
 13. Murk Muller. The RDF Dictionary and the standardisation process in the legal domain. In *Proceedings of XML Europe*, 2002.
 14. H. Pinto, A. Prez, and J. Martins. Some issues on ontology integration, 1999.
 15. 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.
 16. Silvie Spreeuwenberg, Tom van Engers, and Rik Gerrits. The Role of Verification in Improving the Quality of Legal Decisionmaking. In Bart Verheij, Arno Lodder, Ronald Loui, and Antoinette Muntjewerff, editors, *Legal Knowledge and Information Systems (JURIX-2001)*, pages 1–16, Amsterdam, 2001. IOS Press. ISBN 1.58603.201.1.
 17. Andre Valente. *Legal Knowledge Engineering: A Modeling Approach*. PhD thesis, Amsterdam, 1995.
 18. Andre Valente and Joost Breuker. ON-LINE: An architecture for modelling legal information. In *International Conference on Artificial Intelligence and Law (ICAIL-1995)*, pages 307–315, 1995.
 19. C.J.P. van Laer. The Applicability of Comparative Concepts. *European Journal of Comparative Law*, 2.2, 1998.
 20. R.G.F. Winkels, D. Bosscher, A. Boer, and J.A. Breuker. Generating Exception Structures for Legal Information Serving. In Th.F. Gordon, editor, *Proceedings of the Seventh International Conference on Artificial Intelligence and Law (ICAIL-99)*, pages 182–195, New York (NY), 1999. ACM.