

A Study on Translating Regulatory Rules from Natural Language to Defeasible Logic

Adam Wyner¹ and Guido Governatori²

¹ Department of Computing Science, University of Aberdeen, Aberdeen, United Kingdom
azwyner@abdn.ac.uk

² NICTA, Brisbane, Australia
guido.governatori@nicta.com.au

Abstract. Legally binding regulations are expressed in natural language. Yet, we cannot formally or automatically reason with regulations in that form. Defeasible Logic has been used to formally represent the semantic interpretation of regulations; such representations may provide the abstract specification for a machine-readable and processable representation as in LegalRuleML. However, manual translation is prohibitively costly in terms of time, labour, and knowledge. The paper discusses work in progress using the state-of-the-art in automatic translation of a sample of regulatory clauses to a machine readable formal representation and a comparison to correlated Defeasible Logic representations. It outlines some key problems and proposes tasks to address the problems.

1 Introduction

Legal regulations are expressed in natural language.³ To make them automatically processable for reasoning or information extraction, they must be represented in a machine-readable form. There are several approaches to making regulations machine-readable, e.g. linked documents and annotated documents. We focus on the translation of statements in regulations into formal semantic representations that could then be provided to automated deduction engines, which can then be used to check for consistency and redundancy, draw inferences given ground facts, and provide users with meaningful explanations following a consultation, among other processing tasks. The use cases for such translations are very widespread: extracting and formalising relevant rules from regulations to form *rule books* for particular industries; checking for compliance to regulations; serving expert system web-front ends to users, and others.

The language of regulations seems particularly problematic to process. In [1], a range of issues were identified such as the sentence length, clausal embedding, and list structures, which contributed to long parse times or failures to parse. Beyond parsing issues, we want to translate the expressions in regulations into a formal semantic representation to support the sorts of reasoning tasks and use cases mentioned above. Efforts along these lines appear in early work in artificial intelligence and law [2], though without natural language processing (NLP). Some commercial products are available that support aspects of this process and serve the resultant expert systems to users on the web Oracle Policy Management. However, the source material is heavily preprocessed into a controlled language with limited expressivity (on controlled languages, see [3]). An

³ Copyright ©2013 for the individual papers by the papers' authors. Copying permitted only for private and academic purposes. This volume is published and copyrighted by its editors.

open-source, implemented, controlled-language, *Attempto Controlled English* (ACE), has been applied to clinical practice guidelines [4] and to policy-making statements [5] with some, but limited, success. Pilot studies of parsing and semantic representation of regulations with broad coverage, open source tools, C&C/Boxer [6], have been carried out [7]. On the side of logical representations of regulations, there have been efforts to formalise portions of regulation using Defeasible Logic [8]. Machine-readable representations for legal rules, LegalRuleML, have been developed [9].

The studies with ACE and C&C/Boxer highlight two limitations: the output parse and semantic representation given by the tools must be manually checked to accurately correlate to the intended semantic interpretation of the input expression; relatedly, the outputs have not been associated with logical or machine-readable representations that could serve as requirements for the semantic representation. On the other hand, studies using Defeasible Logic and LegalRuleML do not systematically relate to natural language or the issues of acquiring the formal representations from the source material that is represented in natural language. There remains, then, a significant *gap* between natural language source material and formal, machine-processable representations.

In this paper, we discuss a pilot study in which we use C&C/Boxer to translate regulatory statements to semantic representations and then compare the output representations against logical representations in *Defeasible Logic* (DL) that have been *manually* created. By doing so, we gain a better idea of what each form of representation contains, what is gained or lost, how to scope and evaluate such work, the overall process in the analysis, and what next steps are required in order to improve automatic processing of regulatory text.

In Section 2, we provide information about our method, briefly covering the corpora, C&C/Boxer, and DL. A sample of the output from C&C/Boxer applied to the corpora are reported in Section 3. The DL representation of the sentences is presented in Section 4. In Section 5, we discuss the C&C/Boxer and DL representations in comparison as well as future work.

2 Materials and Method

In this section, we present the materials and the method we apply to the materials (e.g. C&C/Boxer and DL).

2.1 Materials

We examine a selection of Section 8.2 of Australia's *Telecommunications Consumer Protections Code* (2012) on complaint management. Broadly speaking, we take a *piece-meal* approach to the overall problem of processing the text, *filtering* and *preprocessing* the original material to some degree to make it *amenable to automatic processing*, yet leaving most of the relevant structure intact. Each of the preprocessing editorial moves is recorded, justified, and systematically applied; however, to economise on space, we suppress discussion of the edits here. The original material contained 173 words, and given the structure of the document, an unclear number of sentences. As the original data has formatting conventions that are not relevant at this point for the semantic content, we have reformatted the data, which we refer to as the *Source Data*. An additional layer of filtering is applied to the *Source Data*, which contains a range of complications which are not relevant to our current exercise such as lists, subordinate clauses, and references. We have manually preprocessed the data, resulting in *Modified Source Data* of 125 words in five sentences:

Modified Source Data

8.2.1.a.xii. Suppliers must advise consumers in everyday language of the resolution of their complaint as soon as practicable after the supplier completes its investigation of the complaint.

8.2.1.a.xiii. A. Suppliers must complete all necessary actions to deliver the resolution offered within 10 working days of the consumer's acceptance of that resolution unless otherwise agreed with the consumer.

8.2.1.a.xiii. B. Suppliers must complete all necessary actions to deliver the resolution offered within 10 working days of the consumer's acceptance of that resolution unless the actions are contingent on actions by the consumer that have not been completed.

8.2.1.a.xi. Suppliers must provide a means for the consumer to monitor the complaint's progress.

8.2.1.a.xiv. Suppliers must only close a complaint with the consent of the consumer or if clause c below has been complied with.

While this is a small corpus, it still allows for instructive semantic representations as well as challenges.

2.2 Method

C&C/Boxer C&C/Boxer automatically parses the sentences of the Modified Source Data and gives an associated semantic representation.⁴ C&C/Boxer consists of a fast, robust *combinatory categorial grammar* (CCG) parser and *Boxer* [6], a tool that provides semantic representations in Discourse Representation Structures (DRSs) of Discourse Representation Theory (DRT) [10] for discourses, including pronominal anaphora and discourse relations. DRSs have equivalent First-order Logic statements in representations that are suitable for FOL theorem provers, e.g. *vampire*.

To economise our presentation, we omit parses and only consider DRSs. We provide a simple illustration of the DRS output for *Bill threw the ball into the street* in Figure 1. In Figure 1, there is a box notation, where boxes represent a knowledge base or

x1 x3 e5 x7 t9 t10
named(x1, bill, nam)
ball(x3)
throw(e5)
Cause(e5, x1)
Theme(e5, x3)
street(x7)
into(e5, x7)
now(t9)
e5 \subseteq t10
t10 < t9

Fig. 1. DRS of *Bill threw the ball into the street*.

Discourse Unit of FOL expressions: a top sub-box represents discourse referents and a lower sub-box the FOL predications. In the example, we have six discourse referents and ten predications (including set and order relations). We discuss some of these. In

⁴ <http://urd.let.rug.nl/basile/gsb/webdemo>

the example, a *named* entity relation is introduced between a variable $x1$, the string *bill*, and the type *nam*. There is an entity which is a ball, another which is a street. Given the neo-Davidsonian, event-theoretic representation [11], we have an event $e5$, which is a throwing event, and thematic roles, one for the *Cause* of the event and another for the *Theme*. Bill is associated with the cause of the throwing and the ball with the theme (the object) that is thrown. Finally, there is temporal information. While there may be some disputes about the semantic representation (e.g. about thematic roles or the interpretation of the preposition), by and large we find this an acceptable semantic representation.

Applying C&C/Boxer to longer, more complex sentences such as in our corpus results in correlatively more complex derivations and semantic representations. Such complex sentences and discourses must be carefully checked that the parse is correct and, more importantly, that the semantic output corresponds to semantic intuitions for an interpretation of the meanings of the sentences (assuming some way to determine these). We illustrate this further later.

Defeasible Logic and Deontic Logic In this section, we give a brief overview of *Defeasible Logic* (DL) [12], which we use to represent rules that are defeasible, and *Deontic Logic*, which represents concepts of *obligation*, *prohibition*, and *permission*.

In the legal domain, rules are well-known to be *non-monotonic*, that is, they admit of exceptions where the rule does not apply or where new information blocks the inference from the rule. DL takes an approach to non-monotonicity that is easy to implement and has been used in various application domains, e.g. regulations, business rules, and contracts [12]. In DL, there are five key features:

- facts - indisputable statements, e.g. *Bill is happy* is $happy(bill)$;
- strict rules - material implication in classical logic, e.g. *Emus are birds* is $r': emu(X) \rightarrow bird(X)$;
- defeasible rules - rules from which we draw inferences, unless the rule is defeated by superior, contrary evidence, e.g. *Birds typically fly* is $r'': bird(X) \Rightarrow fly(X)$;
- defeaters - rules that prevent conclusion of a defeasible rule from holding. They produce contrary evidence, e.g. *If an animal is heavy then it might not be able to fly* is $heavy(X) \rightsquigarrow \neg fly(X)$, which only prevents the conclusion $fly(X)$ where $heavy(X)$;
- a superiority relation among rules - the relation allows us to draw a “winning” conclusion from rules with opposition conclusions, e.g. where $r''': brokenWing(X) \Rightarrow \neg fly(X)$ and $r''' > r''$, the bird with a broken wing does not fly.

A defeasible theory is a program or knowledge base with these features.

In addition to defeasibility, legal reasoning engages the *deontic concepts*, that is, the concepts bearing on *obligation* (**O**), *prohibition* (**PR**), and *permission* (**PER**) along with related concepts of *violation*, where a violation obtains if what is obligated has not been fulfilled or if what is prohibited has come to pass. There area range of subsorts of obligations (see [13] for the subsorts and definitions), where **OM** is relevant to our example:

Maintenance obligation (**OM**) - obligations that, once introduced, require that a state be maintained for a given period of time, e.g. *After opening a bank account, customers must keep a positive balance until bank charges are taken out.*

3 Semantic Representation

In section 2, we presented the corpora and analysis method using C&C/Boxer and Defeasible Logic with deontic operators. C&C/Boxer was applied to the five sentences in our *Modified Source Data*, and every sentence was parsed and given a semantic representation. Essential for our purposes is to consider the semantic representation. In this paper, we only have space to discuss one of the examples.

In Figure 2, we have the representation for statement 8.2.1.a.xi, containing one main DRS, 10 entities, 16 predications, and one subordinated DRS. The main clause *Suppliers must provide a means for the consumer* is paraphrased: the modal *must*, given as \square , has wide scope over the whole representation; the agent of the event of providing is the supplier and the means is the theme; the means are in the *for* relation with the consumer; the time of the event is in the future with respect to *now*. The portion representing the subordinate clause *to monitor the complaint's progress* is paraphrased: a proposition *p2* represents a monitoring event with the supplier as agent, with a progress entity in the *of* relation to a complain, where the progress entity is the location of the monitoring event.

There are several issues to note about the semantic representation. The main clause has an acceptable representation. Semantic operators, e.g. *must* and predications relations, e.g. *for*, are semantically underspecified. Other predications have some intuitive sense, e.g. *Agent*. For the law, some *bearer* of the obligation is required, even if this is universal or generic; in the semantic representation above, there is no such indication of bearer. An important point is that the *generic*, law-like meaning of the sentence, signalled by the plural subject noun in combination with present tense, is not represented. Substantive problems arise with the subordinate clause: *progress* is taken as a location of the event of monitoring rather than a theme, which arises given the lexical specification of the verb *monitor*; the agent of the monitoring is the supplier, rather than the consumer. The first problem relates to the lexical specification of verbs, which are often polysemous. The second problem relates to what is known as *control* such as appears in the difference between the inferred subjects of *leaving* found in *Bill promised Jill to leave* and *Bill persuaded Jill to leave*; there are classes of verbs that behave one way or the other; in the example, *provide* is like *persuade*, not like *promise*. These issues may be resolved through better implementations of thematic role structure and control.

This is an example of the sort of output and analysis available for each of the sentences in our corpus. However, it is difficult to generalise about the outputs or the issues of the semantic representations, as each sentence has particularities that bear further discussion. In the next section, we discuss the related DL representation.

4 Representation in Defeasible Logic

In 2.1, we indicated the *Source Data*, which was used to manually translate into 10 DL rules and one rule ordering. However, as space is limited, we only present the DL representation associated with our C&C/Boxer output and mention aspects of the others.

- Sentence: *Suppliers must provide a means for the consumer to monitor the complaint's progress.*
- DL 8.2.1.a.xi: $\text{complaint}(X), \text{complaint_acknowledgment}(X) \Rightarrow$
[OM]customer_monitor_progress(X)




Fig. 2. 8.2.1.xi: Suppliers must provide a means for the consumer to monitor the complaint's progress.

The method of translation is entirely manual and intuitive. Between the sentence and DL, we see a range of differences: in the DL representation, the *supplier* is missing; complex predicates are introduced into the DL representation that are presuppositions, e.g. *the complaint is acknowledged*; what is linguistically complex is rendered as a DL predicate, e.g. *customer_monitor_progress*; while the deontic operator appears in both the sentence and DL, it appears as complex operator, *maintenance obligation*, which is not clearly associated with the linguistic source (which arises from the generic meaning of the plural subject with present tense); tense is not represented; the bearer of the obligation, e.g. *suppliers*, is not explicit.

5 Discussion

In this section, we discuss observations about the two approaches, how they relate, draw out some general points, and end with future work.

C&C/Boxer automatically provides a parse and a correlated semantic representation for each of our sentences; however, there may be some issues with the accuracy and completeness of the semantic representation. This is, in many respects, an issue to be addressed by refinements to C&C/Boxer itself. The semantic representations are highly articulated, identifying all the individuals, events, and relations, whether found in explicit linguistic forms (e.g. *noun objects*) or implicit (e.g. *thematic roles*). Temporal relations are represented. However, as noted, the generic interpretation is not represented, nor is the bearer of the obligation indicated.

In contrast, DL is a manual translation that represents the meaning of the source clause at a high level of syntactic and semantic abstraction in several respects and in contrast to the C&C/Boxer representations. In DL, complex combinations of words that form a phrase are represented as complex predicates; complex operators such as **OM** are not composed from their parts; temporal order is lost (or subsumed under the interpretation of the defeasible conditional); fine-grained elements of the source material are omitted, e.g. *means*; different participants and their roles in the actions are either omitted or incorporated into a predicate, e.g. *suppliers* are omitted and the customer

appears in *customer_monitor_progress*. The disadvantage of such complex predicates is that syntactic structure and semantic compositionality are largely obscured.

Asides from issues about granularity, the two most significant differences between the C&C/Boxer and DL representations are the representation of defeasibility and the scope of the modal. In the C&C/Boxer examples, no conditional representations arise without explicit, linguistic conditionals (or related operators) in the sentences. Moreover, C&C/Boxer provided only specific rather than generic interpretations, which could be taken to represent defeasibility. In combination, we can say that C&C/Boxer outputs do not represent defeasible rules as in DL. Yet, in natural language semantics, non-monotonicity is usually treated quantificationally, whether with adverbs of quantification [14], as generalised quantifiers [15], or in terms of genericity [16]. The first question is, then, what is the most useful or appropriate representation of defeasibility where we are concerned with the automatic translation from natural language into a formal representation? A second related question is to what extent can a tool such as C&C/Boxer accommodate the chosen representation or, turning it around, to what extent ought DL be revised to accommodate natural language semantics of non-monotonicity? Turning to the modal, the scope with respect to a conditional is a complex, largely unresolved matter in natural language semantics [17].

Some of the differences outline above may be taken informatively, in the sense that they indicate how each approach might incorporate or adapt to useful components of the other. There may be ways to bridge the differences; for example, complex predicates can be systematically related to component parts, and quantificational representations of non-monotonic operators could be translated into correlated statements of defeasible logic. But, bridging the differences requires first identifying what those are and whether to bridge them.

We have not discussed evaluation. In statistical or machine-learning analyses, results are usually provided in terms of precision and recall measures, where the performance of a proposed algorithm is measured against a *gold standard* corpus. However, in the absence of such gold standards, we cannot provide such measures; and the creation of such corpora rest on the *specification of what the corpora ought to encode*, which in our view, remains unclear in the research community. Rather, the results reported here bear on: (1) the extent to which existing technologies produce more or less intuitively accurate output; and (2) specific observations about the outputs in comparison; and (3), setting an agenda for future research. Nonetheless, for future work, some explicit measures for evaluation of each approach must be provided. This is tied to the issue of requirements; while initially it seemed that DL representations could be used as abstract specifications to which C&C/Boxer should fulfill, this is not clear. Indeed, this study only serves to highlight that the two approaches have rather different means and objectives, even if somewhat related. However, these topics must be for future work.

References

1. Wyner, A., Peters, W.: On rule extraction from regulations. In Atkinson, K., ed.: Proceedings of the 24th International Conference on Legal Knowledge and Information Systems (JURIX 2011), Vienna, IOS Press (2011) 113–122

2. Sergot, M., Sadri, F., Kowalski, R., Kriwaczek, F., Hammond, P., Cory, T.: The British Nationality Act as a logic program. *Communications of the ACM* **29**(5) (1986) 370–386
3. Wyner, A., Angelov, K., Barzdins, G., Damljanovic, D., Davis, B., Fuchs, N., Hoefler, S., Jones, K., Kaljurand, K., Kuhn, T., Luts, M., Pool, J., Rosner, M., Schwitter, R., Sowa, J.: On controlled natural languages: properties and prospects. In: *Proceedings of the 2009 conference on Controlled natural language. CNL'09, Berlin, Heidelberg, Springer-Verlag* (2010) 281–289
4. Shiffman, R.N., Michel, G., Krauthammer, M., Fuchs, N.E., Kaljurand, K., Kuhn, T.: Writing clinical practice guidelines in controlled natural language. In: *Proceedings of the 2009 conference on Controlled natural language. CNL'09, Berlin, Heidelberg, Springer-Verlag* (2010) 265–280
5. Wyner, A., van Engers, T., Bahreini, K.: From policy-making statements to first-order logic. In: *Proceedings of International Conference on Electronic Government and the Information Systems Perspective EGOVIS-2010*. (2010) 47–61
6. Bos, J.: Wide-coverage semantic analysis with boxer. In Bos, J., Delmonte, R., eds.: *Semantics in Text Processing. STEP 2008 Conference Proceedings*. *Research in Computational Semantics, College Publications* (2008) 277–286
7. Wyner, A., Bos, J., Basile, V., Quaresma, P.: An empirical approach to the semantic representation of law. In: *Proceedings of 25th International Conference on Legal Knowledge and Information Systems (JURIX 2012)*, Amsterdam, IOS Press (2012) 177–180
8. Antoniou, G., Billington, D., Governatori, G., Maher, M.J.: On the modeling and analysis of regulations. In: *Proceedings of the Australian Conference Information Systems*. (1999) 20–29
9. Athan, T., Boley, H., Governatori, G., Palmirani, M., Paschke, A., Wyner, A.: OASIS Legal-RuleML. In: *Proceedings of the 14th International Conference on Artificial Intelligence and Law (ICAIL 2013)*, Rome, ACM (2013) xx–xx To appear.
10. Kamp, H., Reyle, U.: *From Discourse to Logic: Introduction to Model-theoretic Semantics of Natural Language: Formal Logic and Discourse Representation Theory*. Springer (1993)
11. Parsons, T.: *Events in the Semantics of English: a Study in Subatomic Semantics*. MIT Press (1990)
12. Antoniou, G., Billington, D., Governatori, G., Maher, M.J.: Representation results for defeasible logic. *ACM Trans. Comput. Log.* **2**(2) (2001) 255–287
13. Governatori, G., Rotolo, A.: Norm compliance in business process modeling. In: *Proceedings of the 2010 international conference on Semantic web rules. RuleML'10, Berlin, Heidelberg, Springer-Verlag* (2010) 194–209
14. Lewis, D.: Adverbs of quantification. In: *Formal Semantics of Natural Language*. Cambridge University Press (1975) 178–188
15. Barwise, J., Cooper, R.: Generalized quantifiers and natural language. *Linguistics and Philosophy* **4** (1981) 159–219
16. Pelletier, F., Carlson, G.: *The Generic Book*. The University of Chicago Press (1995)
17. Wyner, A.Z.: *Violations and Fulfillments in the Formal Representation of Contracts*. PhD thesis, Department of Computer Science, King's College London (2008)