



UvA-DARE (Digital Academic Repository)

Legal theory, sources of law and the semantic web

Boer, A.

DOI

[10.3233/978-1-60750-003-2-i](https://doi.org/10.3233/978-1-60750-003-2-i)

Publication date

2009

[Link to publication](#)

Citation for published version (APA):

Boer, A. (2009). *Legal theory, sources of law and the semantic web*. [Thesis, fully internal, Universiteit van Amsterdam]. IOS Press. <https://doi.org/10.3233/978-1-60750-003-2-i>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 7

Summary and Conclusions

7.1. Introduction

This book aims to be a work of ontology: an account of relevant aspects of the knowledge domain of law from the perspective of a legal knowledge engineer interested in sources of law. One cannot however say that the result of this work *is* an ontology; This book presents a mix of design principles, design patterns for knowledge representation in OWL DL, and ontology fragments.

Section 7.2 summarizes the essentials of the conceptual framework constructed throughout this book.

MetaLex and the Legal Knowledge Interchange Format (LKIF) ontology played an important background role. The notions of representation and reference (to concepts through terms) were introduced by me in MetaLex, and betterness in relation to normative rules was introduced by me in the LKIF ontology. Three concepts that play a central role in this book are *not* part of either ontology: the notion of *constitutiveness*, the notion of *applicability*, and the notion of *intentionality* as the *execution* of a plan. Also the notion of playing an *agent role* is not found there.

Section 7.3 focuses on a number of conclusions, both positive and negative, that can be drawn from the experience of constructing this conceptual framework. These conclusions address the questions of section 1.3. Section 7.4, which ends this exposition, addresses a number of interesting avenues for future research that are beyond the scope of this work as outlined in chapter 1.

7.2. Summary

Based on experiences over the last years with the MetaLex standardization effort and with the Legal Knowledge Interchange Format, this book discusses a number of relevant features of legal reasoning, and approaches for isolating different types of legal knowledge.

The legal knowledge source par excellence is the formal (written) *source of law*, discussed in chapter 5. The formal source of law is a bibliographic entity: a document. It exists in different ontological strata.

It is initially a physical *item*, that can be copied into additional items, which are however still examples of the same document *manifestation*. It can be edited, without changing the information represented by the document, producing additional manifestations of the same document *expression*. Even the information contents of the document can be changed, producing alternative expressions, while it still remains the same *work* from the perspective of bibliographic identification. The work level is only relevant if there is common content in all expressions of the work, and if later expressions share bibliographic identifying data with the original expression.

The XML standards on which the Semantic Web is built tend to identify documents on the item and manifestation level, while the legal field identifies documents on the expression (consolidated versions) and work level. For knowledge representation the expression level is the level of interest, as it is the level at which documents are identified and individuated by their information contents.

Chapter 4 described law in terms of institutions whose primary purpose is to create normative order by way of formalization – the sources of law specify a *formal, institutionalized normative order*. An *institution* is a *social* artifact: the structures of the institution are defined by the institutional facts that make up the institution, and its mechanisms of change are the constitutive rules that specify what *constitutes*, or counts as, an institutional fact.

The source of law is a writing that can be, is, was, or presumably will be used to back an argument concerning the existence of a certain institutional entity in a certain legal institution: The source of law is the result of a legislative act performed with the intent of creating that institutional entity, and functions as evidence of the legislative act.

The *legislative act* belongs to a broader category of *legal acts* that are characterized by 1) the requirement that one *intends* to bring about a certain institutional change, and 2) that this intent is communicated in writing. The legislative act distinguishes itself by the nature of the institutional entities it creates: legal rules. Other legally relevant acts also change legal institutional reality, but do not create recognized legal rules. Legislation may also achieve some other auxiliary effects, relevant to the operation of legal rules. The opinion that sources of law represent legal rules is so widely shared that it was uncritically adopted in this work.

The main function of the source of law is to function as evidence that certain institutional events, which changed the state the institution is in, happened. The sources of law function as a log book of relevant legislative changes to a legal institution.

A well-picked body of sources of law may also be considered a blueprint or a snapshot of the rules and structures of a specific legal institution of interest in a certain time frame, if the legislator does his work as a designer well. The state of the documentation and the state of the institution do not however necessarily simulate the other in either direction, as section 5.2 showed, even though the information contained in the documentation is essential for the social recognition of the institution. The sources of law and the institution, and the legal rule and the specific text fragment used as evidence for its existence, should be distinguished to avoid unnecessary complications in version management.

The task of the legal knowledge engineer, who is interested in representing legal knowledge, is primarily to reconstruct and represent the rules and structures of the institution rather than the source of law, which is itself a representation of (changes to) the rules and structures of the institution.

7.2.1. Reference and Representation

A specific goal of the recommendations in this book is liberating knowledge engineers from the self-imposed straightjacket of trying to represent a legal rule by a logical proposition. The logical propositions are statements *about* the legal rules and other relevant entities in the institution. A good argument against the bijective mapping of legal rules to logical propositions was found in [195] (in the context of norms; cf. section 2.2.2): legal rules have no truth value and truth-functional connectives should not be applied to them.

Representation of a source of law expression thus consists of two initial interpretative phases, although it doesn't end there:

1. Identifying the entities in institutional reality *first* represented by the source of law, as these were initiated by the legislative act evidenced by the source of law; and
2. formulating the logical propositions that constrain the valid models of institutional reality and describe the interfaces by which institutional reality can be changed, as evidenced by the sources of law and the way they are used.

The knowledge representation describes institutional entities as statives – entities whose existence is limited in time – linked to the sources as evidence of their existence through the mechanisms of *reference* and *representation*.

Section 5.4 proposed a number of principles for identification of institutional entities in sources of law. The source of law at the expression level is identified by a URI, *refers* to a set of URI-identified *terms*, and it *represents* a set of URI-identified *legal rules* and other *legal assertions*. This applies to the source of law as a whole, and to any URI-identifiable fragment of the expression. Since the expression does not change – it is only created – one would expect that the set of legal rules it represents and the set of terms it refers to also normally do not change. In the rare case that they do, we are dealing with an *ex tunc* expression, which is the result of a legal fiction.

Terms can be categorized into those that are *in* the institutional reality created or modified by the legal rules, and those that are relevant to this institutional reality but outside the boundaries of the institution. Legally relevant terms are not defined by the institution: whether the term does or does not accurately describe a certain entity is from the point of view of the institution a *technical question* (cf. section 2.5).

All legal terms and rules *first* identified in an expression-level source of law should be assigned the same namespace, which is bijectively related to that source of law. No other terms are in that namespace. A new URI, in a new namespace, can be assigned to a rule or term only if a rule or concept is deemed to change in a new expression of a work, in the sense that the old concept and the new concept denoted by the same term in the source cannot be reconciled. The decision to integrate terminology from different expressions is made explicit in the form of equivalence statements.

One might think of the set of terms and legal rules occurring in all known expressions of a work as the shared set of terms and rules at the work level, but the shared work level set only exists from a specific vantage point in time, or only once the source of law has become immutable after its repeal, and the shared work level set – which can no longer change – has become largely irrelevant. This abstraction cannot be safely used, since it is likely to lead to versioning problems.

The source of law on the expression level however *cites* other rules on the work level, while the legal rules we represent are necessarily identified by their *first representation* in an expression, and not on the work level. A citation (*text fragment*) *w* *applies to* (*concept*) *C* should for instance be read as *each legal rule that is represented by an expression-level text fragment that embodies work fragment w applies to C*.

7.2.2. Legal Rules and Logical Propositions

Relevant patterns in logical propositions describing the functions of legal rules revolve around the notions of *constitutiveness* and *applicability*. In particular the notion of constitutiveness, introduced in section 4.2, is central to the institutional interpretation of law.

The legal rules represented by the source of law appeal to two separate realities – institutional reality and brute reality – and perform a mapping from brute reality – the ontological substratum – into institutional reality – the ontological superstratum. The substratum has an existence independent of the rules, while the superstratum is supervenient on the substratum and exists by virtue of social recognition of the rules of the institution. Institutional events are *constituted by* events in brute reality.

Institutional rules map out a logical space of possible models of the institution: they form the institution’s ontology, and can be interpreted as terminological axioms. They are not defeasible as a matter of convention: if the suspicion arises that a certain institution has an ambiguous internal structure, this should serve as a hint to either redefine the boundary between the institution and brute reality, or to split the institution into two or more institutions.

Since the quality of the mapping between brute and institutional realities is likely to be imperfect, we should assume – as a matter of convention – that institutional facts only exist as long as it is consistent to believe they exist.

The main function of the *constitutive rule* is to define the interface through which the state of the institution can be changed. As such, the most important question they answer is how to map brute reality *into* institutional reality. Legal rules of a constitutive nature are however not conditionals which take brute reality propositions as a premise, and institutional propositions as conclusion. The constitutive legal rule may have a variety of functions, corresponding to certain patterns of logical propositions describing them.

Logical propositions *describing* constitutive legal rules (cf. sections 4.4.1 and 5.2) can be categorized into *requirements*, which posit necessary conditions on the mapping between brute and institutional fact, and *indicators*, which license an inference from brute fact to institutional fact, if and as long as it is consistent to do so. In addition, the law may impose constraints on knowledge availability: there must be positive justification from brute reality – not from the legal rules alone – that a change to institutional reality happened. One cannot freely infer institutional facts from lack of knowledge. These rules were tentatively called *burden of proof* rules. In an argumentation context, these function as (epistemic) obligations to justify arguments with certain information that must be collected (usually by questioning a user).

Requirements, and the terminological axioms used to describe institutional rules, can be accurately captured by straightforward OWL axioms. Indicators and burden of proof rules are however autoepistemic in character.

Applicability, mainly discussed in section 5.2.2 and 5.2.3, plays a central role as soon as the logical proposition and the legal rule are separated. If a logical proposition states that a certain proposition indicates or entails another proposition, we must add to this conclusion that a legal rule has been applied. Any logical proposition about

the application of a legal rule must conclude that the legal rule has been applied.

A special class of legal rules (generally described as requirements) constrains the applicability of other rules. *Applicability rules* are used to avoid repetition of the same requirement, to demarcate the extent of jurisdiction over people, territory, and substance, and applicability in time, and to restrict the applicability of legal rules made by a lower legislator. The *choice rule* makes the application of one legal rule conditional on the application of another legal rule, if the combined application of two legal rules is judged to be logically ambiguous or *contrary to the intended normative order* (cf. section 6.6).

Applicability rules can be used as a guide for clustering sources of law into coherent domains of application, and are generally used for that purpose by the law. They should also play that role in legal knowledge engineering, but generally do not do so explicitly. This book explains legal fiction, and the closely related also-applicability pattern, in terms of applicability: legal fictions are inconsistent with the conscientious use of applicability rules (cf. sections 5.2.4 and 5.4.3).

Sections 4.6 and 6.2 also distinguish the *normative rule*: it is a constitutive rule characterized by the type of institutional fact it establishes and the intention with which it is created. Its logical pattern does not *have* to be distinguished from the already established forms if one doesn't attempt to represent the *intended normative order*.

This observation sets this book apart from previous uses of Searle's constitutive rule (cf. [249, 248]) in legal knowledge engineering. Although Searle subscribes to the logical distinction between constitutive rules and "regulatory" or normative rules, the *legal* normative rule should clearly be considered a constitutive rule (cf. section 4.6). Treating only *some* legal rules as constitutive only confuses the issue.

Note that the logical propositions following the identified patterns are not found in the source of law: they are created by the knowledge engineer. The logic of grouping URI identifying logical propositions into namespaces, and the sentences themselves into files, and doing version management on those files, is a knowledge engineering concern not constrained by the design rationale of the sources of law. The representation proposed in chapter 5 can be expected to be compatible with other methodological concerns besides isomorphism with the source.

7.2.3. Normative Order

The work of the legal knowledge engineer is not limited to reconstructing and representing the rules and structures of the institution. One could argue that the task of creating a reusable knowledge component out of a source of law ends there, but the resulting knowledge components have little purpose if not augmented with another interpretation at yet another level of existence: the *intended normative order* (cf. chapter 6).

An institution can intend and attempt to formalize a normative order, but the existence of normative order itself is only proven by regularities in the behaviour of individuals in a community.

Normative order always exists, regardless of whether it is formalized, and the

legislator usually just tinkers with the details of an existing normative order in society. Sometimes the actions of the legislator only confirm, or perform minor repairs in existing, spontaneously arising normative order, and sometimes the legislator only tries to create the right conditions so that acting in one's own interest will create the intended effects.

Even for the mere purpose of complying with the law, we have to ascribe to the legislator the intent to formalize a specific normative order. The *normative rule*, discussed in section 6.2, has a straightforward translation: what it *tells* us is that a certain behaviour brings about the institutional fact of violation of the normative rule, and what it *means* in terms of the intended normative order is that the legislator intends us to choose against execution of plans that could be recognized by others as a violation of the normative rule: one could call this the *deontic choice* effect of the rule.

Observe that deontic logics represent not the legal rules, but the intended normative order as readily apparent from the normative rules. As observed in sections 4.6 and 6.2 a straightforward assessment framework classifying cases as allowed or disallowed suffices for the institutional meaning of the normative rules of the institution. Since the sources of law do not formalize the whole intended normative order, a representation of these rules in deontic logic gives us only a partial account of behavioural regularities in a community.

The legislator acts on an existing normative order, and performs legislative actions with the intent to change it. The legislator relies on three important mechanisms to effect changes in the normative order: public announcement of the rules and structures of the legal institution and the competence of addressees of his announcements to understand the intended implications for the normative order, the widely held preference for predictable social interactions among the addressees, and the ability of the widely recognized legal institution to effectively organize punishment of those that don't comply with the prevailing normative order, which underlies the deontic choice effect of the normative rule.

That the ascribed intended normative order – the whole of it – has direct relevance for the application of normative rules becomes clear in situations of normative conflict and compliance dilemma: if compliance with the normative rules fails to achieve the intended normative order, this calls into question their *applicability* to the case even if there is no plausible argument that application of the rule would create an inconsistent institutional reality (cf. section 6.2, specifically 6.3.1).

Just like the sources of law cannot be taken at face value as a blueprint of a legal institution, the normative reading of its rules cannot be taken at face value as a blueprint of the normative order it formalizes.

This is the most important reason why in the field of comparative law the importance of taking into account the cultural and social context in which legislation exists is often stressed, and why *legal transplantation* of legislation doesn't always have the expected effects (cf. section 6.7.3). Filling in the gaps is an interpretative task, and involves ascribing a clear intent to collective agents – parliaments, governments – that very likely had no such monolithic intent in the first place.

In many cases presuppositions about the pre-existing normative order being re-

gulated are found directly in the source of law: Hohfeldian transactional notions like rights correlative to a duty, liability correlative to a power, etc, are founded in certain expectations associated with the adoption of certain agent roles in existing social scripts (cf. sections 4.8.2 and 6.7.1). The legislator may for instance choose to phrase an obligation in terms of a right of a counterparty; This does not change the fact that the rule should be interpreted primarily as an indirect obligation, but it does yield information about the intended normative order. Another clear source of presuppositions about the intended normative order is found in the evaluative nature of the terms used: legal qualifications are often either value-positive or value-negative (cf. section 6.7.1).

Legislative explanatory documents accompanying legislation also often explain the intended normative order. Articulations of the intended normative order play a direct role in the form of *norms of analysis* in the legislative drafting phase (cf. section 6.7.3). As argued in section 6.7.2, courts also ascribe an intended normative order, and latch on to deficiencies in the operation of the rules relative to this ascribed normative order to create legal rules of their own: applicability rules which make the application of one rule conditional on the application of another rule if certain conditions are met.

Consider also the perceived problem of giving John permission to live in a certain house and Jane permission to destroy it, discussed in section 6.3. It is a problem only because we immediately ascribe to John a preference for i.a. staying alive, and to the legislator awareness of John's preference and the willingness of taking into account John's preference for staying alive: the legislator is therefore behaving irrationally in our opinion in his attempts to create normative order.

As observed in section 6.7.3, analysis of comparative arguments about differences and similarities between legal regimes has the potential to teach us the difference between the institution and the normative order it proposes, and requires a great degree of intercoder reliability in the representation of normative order: legal knowledge engineers can use this information for more robust legal knowledge representation. At the same we have to recognize that there is little demand for LKBS that evaluate proposed legislation, and little prospect for investment in the development of such LKBS. Comparison for the purpose of assisting the mobility of people, goods, and services in the EU, or for preparing the implementation of legislative change in administrative settings *is* on the other hand of recognized relevance to society.

7.2.4. Subjunctive Betterness

In the representation of normative order the *subjunctive betterness* conditional plays a central role in this book (cf. section 4.6 and chapter 6 generally, section 6.2 specifically). In OWL, the legal normative rule of obligatory form $O_n(\alpha \mid \beta)$, where n is the reified normative legal rule, can be represented as follows:

$$\begin{aligned}\alpha \sqcap \beta &\sqsubseteq \forall \preceq_n \alpha \sqcup \neg\beta \\ \neg\alpha \sqcap \beta &\sqsubseteq \exists \preceq_n \alpha \sqcap \beta\end{aligned}$$

If the obligation is for instance *when driving a vehicle one should keep to the right of the road*, the interpretation in terms of subjunctive betterness is:

When driving a vehicle and keeping to the right of the road, all equal or better alternatives would involve not driving a vehicle or keeping to the right of the road. When driving a vehicle and not keeping to the right of the road, an equal or better alternative would be driving a vehicle and keeping to the right of the road.

The combination of both statements is only satisfied by considering alternatives involving not keeping to the right while driving strictly worse than those that do involve keeping to the right, or by not considering alternatives which involve driving a vehicle at all.

Preferences – whether communal or private – can be represented in the same way (cf. section 6.7.1). Normative order arises from the fact that a large body of such statements – representing legal, communal, and private normative rules – are generally accurate descriptions of the behaviour of a community because people tend to copy both the behaviour and the justification of that behaviour of those around them.

This book also sketches how this information about the prevailing normative order is organized and applied (cf. section 4.5): Kahnemann's reference transaction is a script for getting something done using other people as a resource. The script generally works because the other participants act in accordance with their agent role in the script: they understand and subscribe to the prevailing normative order. The *agent role* functions as an anchor for ascribed preferences: someone is not considered a *seller of snow shovels* because one merely has the competence or ability to do so, but because others expect that they can engage with this seller in a predictable transaction for obtaining a snow shovel in exchange for some predictable amount of money.

Unfairness is deviation from the scripted execution of a transaction to the advantage of another participant and to the disadvantage of you, without an acceptable justification.

As long as we accept that a large body of such subjunctive betterness propositions are shared, and that propositions like *I prefer that you . . .* are irrelevant to describing a prevailing normative order, subjunctive betterness propositions generally do not have to be indexed to their origin (i.e. we do not have to distinguish between those in your and those in my mind). Even the ones of legal origin do not have to be distinguished as to their origin, as long as their context of applicability is suitably constrained to time period and jurisdiction by applicability rules.

7.3. Conclusions

For the purposes of maintenance of legal knowledge-based systems, a one-to-one mapping between knowledge components and sources of law, and between fragments of text in the source of law and logical propositions, would superficially appear to be ideal. This expectation is based on the assumption that the legislator's design rationale for sources of law and the knowledge engineer's design rationale for knowledge components can be functionally aligned. The Semantic Web vision moreover suggests that such knowledge components are reusable, or can be designed for reusability.

The problem of making high-quality reusable knowledge components representing sources of law available on the Semantic Web can be approached from several angles, as indicated in the introduction of this book.

7.3.1. The Semantic Web

In section 1.3 a technical theme was addressed: how does one integrate existing core technologies and technical standards of the Semantic Web, like its naming and addressing standards, and the description logic OWL DL, with legal knowledge representation?

Of central importance is the uniform resource identifier or URI, and the realization that the URI violates the unique name assumption typical in knowledge representation. This has consequences for inference in OWL DL, as explained in section 3.4.

OWL DL is a monotonic description logic, designed specifically for robustness in a distributed and fault-tolerant environment. Description logic is particularly useful for describing terminological axioms that are true by definition. It however lacks the possibility of expressing propositions about propositions, and does not support non-classical, but for KBS development important, features like defaults, integrity constraints, and negation as failure. These are certainly necessary for intelligent behaviour, as section 2.4.1 explained. Arbitrary use of such features would however certainly not help to achieve our goal of minimizing impact in the deployment of updated knowledge components.

Chapter 3 sketched a general layout of straightforward knowledge-based systems using OWL knowledge components, that accounts for the controlled use of such constructs. A distinction was introduced between a number of types of logical proposition: OWL DL's terminological axioms and assertions, on the one hand, and default rules and integrity constraints on the other hand. This distinction describes several existing methods for interfacing KBS to OWL DL. Moreover, many proposals to extend OWL DL focus on such constructs.

The very general level of abstraction at which these constructs were described, permitted description of the relation of such constructs to particular types of legal rule without undue epistemological commitments to a specific problem solving strategy.

7.3.2. Ontological Stratification

The great challenge is to structure legal knowledge in accordance with the distinction made between types of logical propositions, and into coherent, justifiable, and reusable, knowledge components. Chapter 2 introduced some theoretical devices for this purpose that classify knowledge by function. To give a philosophical justification for localizing the use of default rules and similar defeasible constructs, a number of concepts were introduced: abstraction, aggregation and refinement of abstractions, and ontological stratification (cf. section 2.3.1).

To effectively apply these theoretical devices to the domain of law, and specifically sources of law, the inspiration must be found in legal theory (generally chapter 4).

Knowledge in its traditional sense, as attainment of truth, usually describes “islands” of abstract knowledge that can be considered axiomatic: *if* an abstract theory is applicable, it is indefeasibly applicable as a whole, but *whether it is advisable to try* depends on the thing one want to apply the theory to, and on the problem one is solving. Abstract theories always come with if's and but's. Terminological axi-

oms accurately model these theories, while default rules function as indicator that the abstract theory *may* be applicable.

This philosophical justification accounts for defeasible rules, for problem solving techniques like heuristic classification, etc. These are alternative epistemological accounts of the same observations about the character of reasoning with defeasible knowledge.

The nasty details are in the abstraction moves between abstract theories: an abstraction is only permissible as long as it remains consistent with the current belief set, and only useful depending on the problem one is solving. It is these abstract theories whose terminological core can be expected to remain stable over long time periods. These “islands” of stable terminological knowledge, surrounded by rules recommending a jump to the island, are what is captured by a reusable knowledge component.

The concept of an ontological stratum suggests one should look for ontological criteria to tell apart the abstraction and the thing it is an abstraction of. This book proposes several abstractions between ontological strata relevant to law, based on an analysis of central concepts in legal theory relevant to legal knowledge engineers, notably constitutiveness, execution, and normativity (generally chapters 4 and 6), and proposes that defeasibility is limited to such abstractions.

7.3.3. Reusable Knowledge Components

When designing standardized components it obviously makes sense to focus on the functions of the component; These characterize how one interacts with the component. These characterizations of function can be identified at different levels of abstraction, as indicated, and depend on the kind of purpose one has in mind. This of course also applies to knowledge components, and to the source of law.

In this book we have focused on different levels at which we can characterize the source of law (generally chapters 4 and 6):

1. as an account, physical evidence, of legislative changes to legal institutions, and
2. therefore as an indirect and incomplete representation of the structures and the interface of legal institutions, or
3. as the physical product of acts intended to change prevailing normative order towards an intended normative order, and
4. therefore as an indirect, incomplete, and approximate representation of normative order.

This book concludes that constitutiveness and normativity are *different accounts* of the functions of the source of law, *not* alternative functions of legal rules as Searle’s distinction between constitutive and regulative rules, taken up in legal knowledge representation literature, appears to suggest.

The obvious choice is to characterize the function of the source of law, and therefore of the knowledge component, in terms of the legal institution(s) it provides knowledge of. This for instance means that even the normative rule must be in first instance accounted for as a constitutive rule.

The relation between the source of law and the institution is clearly a more direct one than the relation between the source of law and normative order; Effects on

normative order are rather an effect of the legal institution and its social recognition than of the source of law *per se*.

A second reason to discount the relation between source of law and normative order is that law only changes a pervasive, existing normative order, and not always in predictable or well-understood ways.

A third reason to discount the strength of this relation is that even the *intended* normative order, the thing the legislator's actions are aimed at, is at best something to be found "between the lines" in the source of law, if the ascription of such an intended coherent normative order is an accurate model of legislating at all.

The notion that the sources of law give a full account of the structures and rules of the institution is only tenable to the extent that the legislator provides full evidence of it. Since the legislator does have good reason to provide such an account, this concept will work reasonably well. There are however occasional counter-examples, as indicated in chapter 5.

7.3.4. Isomorphism with Sources of Law

The relation between the source of law and the knowledge component is of course easier to establish than the relation between a knowledge component and an institution. The source of law may account for multiple institutions, and generally does so only partially.

The traceable relation between the source of law and the knowledge representation is not found in a correspondence between logical propositions and fragments of the text. This position is suspect philosophically *if* it means that these logical propositions must be connected by truth-functional connectives (which is by the way not the only account one could give of them), and – more importantly – it is highly impractical in a language like OWL DL that limits propositions about propositions.

The position taken in this book is that the source of law *refers* to concepts (through the use of terms) and *represents* legal things, amongst others legal rules. Legislative acts create, change, or destroy these things, and in addition produce physical evidence and representation, in the form of the text of the source of law, of these creations, changes, and destructions. The text is constitutive of the legal thing.

Adaptation of the knowledge representation to changes of the underlying sources of law relates to change of these (institutional) concepts and the legal rules, and not to logical propositions.

Both terms (or concept names), and individual objects are uniquely (but not exclusively) identified by a URI in OWL DL axioms. Although logical propositions about these named entities *can be* identified by URI in OWL DL, this URI is not accessible for reasoning within the *same* OWL DL theory. Legal rules will however directly refer to text fragments in sources of law, in order to refer to other legal rules. Section 5.4.1 gave an example of how this is addressed within the same OWL DL theory.

Traceability to logical propositions is provided by the principle that the logical proposition is *about* the application of a legal rule to some thing (through the **applicable** property).

7.3.5. Legal and Legally Relevant

The analysis of sources of law in terms of institutions supplies us with a criterium to tell apart legal knowledge from legally relevant knowledge, which was one of the objectives set for this work in section 1.3: legal knowledge concerns knowledge of institutional concepts, institutional facts, and institutional rules, and the constitutive rules that tell us how brute facts are abstracted into institutional ones.

Legally relevant are the concepts that describe the brute reality that is abstracted into legal institutional reality by constitutive rules, and – arguably – various epistemological theories that help explain how law enters the picture in planning and plan recognition, how planning dilemma’s should be resolved etc. Legally relevant knowledge is not found in the sources of law, and therefore has no place in the knowledge components of interest in this book.

Constitutive rules do not only indicate, but also restrict, the possibility of abstraction into institutional reality. This is why a distinction between indicators and requirements was introduced. Constitutiveness is based on an ontological criterium: if the antecedent and the consequent of the rule are in different strata, the rule is constitutive.

This expands the notion of constitutive rule beyond the notion of *supervenience conditionals*, found in literature, which only explains indicators.

It has been claimed in legal knowledge engineering (cf. for instance [130], and the references given there in the appropriate section) that one of the strange features of constitutive legal rules and legal facts is that a legal fact can only come into existence if one of the constitutive legal rules that has it as a conclusion is applicable; This would appear to mean that, as eloquently formulated in [130]:

[..] two or more supervenience conditionals cannot all be seen as definitions at the same time. They rather seem to imply, together, a definition having as antecedent the disjunction of the antecedents of the different conditionals. If consenting supervenes on signing a document, and it supervenes also on signing via internet an e-form but on nothing else, then signing a document or e-signing imply consenting and conversely consenting implies signing or e-signing.

This clearly only explains something about indicators. Moreover the notion of a disjunction of the antecedents of the various conditionals is inherently problematic for maintenance and should be avoided.

This interesting feature of constitutive rules has not been reconstructed in this book through the use of logical form, but of an ontological commitment: each institutional thing is constituted by a constituting thing. In the case of the law, each legal thing is constituted by a legally relevant thing, to which a legal rule has been applied. That this mapping from constituting thing to institutional thing is only constructed through the use of the rules of the institution is a matter of social recognition: the rules of the institution specify its recognized interface.

[130]’s observation applies equally to for instance the classification of diseases; Only certain symptoms are recognized as indicators of a disease. In this respect the interface of the institution and for instance the interface of a component or for instance the user interface of a computer application are also similar: the interface specification is an

exhaustive description of the recognized – i.e. intended and expected – interactions you can have with it. This is also the case with legal institutions and their public interface.

Because the institution however only exists *because* of social recognition alternative interactions would only arise through social recognition of those interactions, while for instance a computer can be directly interacted with in unintended and unexpected ways with for instance a hammer.

7.3.6. Loose Ends

There are several loose ends in this proposal that may cause some uneasiness. Most important of these are the notion of boundaries between legal institutions, and the positioning of normative order.

The assumption that institutional reality *is* internally coherent, as opposed to the defeasible rule interfacing institutional reality and external reality, creates a new problem of identifying boundaries between different legal institutional realities.

External institutional realities that are recognized as legally relevant by a legal institution complicate matters, as these external institutions compete for normative order, may use the same vocabulary, and generally confuse matters for people that recognize the rules of multiple legal institutions and have trouble telling them apart.

The proposals in this book replace the problem of keeping track of defeasibility between potentially all legal rules with the rather arbitrary notion of splitting the law up into as many different legal institutions as required for delegating defeasibility to the interfaces between them. We for instance let go of the idea that one legislator manages one institutional reality: a legislator may unintentionally split one institutional reality into two by mere use of legislative technique.

The problem of telling apart institutions does not however appear to be *more* complicated than determining whether a rule may be defeated by another one. It supplies a new minimality criterium for defeasible reasoning in law.

Lastly, if legal knowledge is in essence limited to knowledge of the institution, then normativity is a mere derivative product of it. The rules of the institution give rise to preferences and to expectations about the behaviour of other people. This is not a special observation about legal knowledge: all knowledge has the purpose of giving rise to preferences and expectations, and all knowledge about social conventions and institutions give rise to preferences and to expectations about the behaviour of other people. There is no a priori reason why there should be a specifically legal account of this.

In law the applicability of legal rules may however depend on normative conflict, which can only be explained in terms of an account of normative order. Also the comparison between *similar* legal institutions, for instance from different countries is an ill-defined question if it does not include the comparison of two intended and prevailing normative orders. There are multiple ways to create the same normative order through different institutional rules and structures, if we take into account that the institution only tries to *change* an existing, prevailing normative order.

In chapter 6 a not entirely successful attempt was made to reconstruct the ma-

major categories of normative conflict in literature based on a simple representation of preferences in OWL DL. An abstraction to preference, or for instance to obligation, attributed to agent roles is justifiable for normative rules if we assume they are constitutive.

Agent roles play an important role in organizing out knowledge about behaviour, as pointed out in chapter 4: when one assumes an agent role, one assumes the preferences associated to it, and when one attributes an agent role to another one has expectations based on these exact same preferences as pointed out in chapter 6. It makes sense to assume that the preferences associated to an agent role should form a coherent whole, but a person may of course have to deal with the expectations based on different attributed agent roles.

Both the comparison of legal institutions and the resolution of normative conflict is hard to conceptualize without the context of specific scripts and agent roles, as chapters 4 and 6 pointed out.

While a construction of the intended normative order, in the context of the relevant scripts, is necessary to realize the wonderful promises of section 1.2, this construction cannot however be considered to be part of the knowledge components representing sources of law. This for instance also means that Hohfeldian concepts like right, liability, etc. cannot be completely accounted for; Whether one “has a right” depends on preferences and expectations, and not on a specific phrasing of a legal rule in a source of law.

This leads to an important conclusion that affects reusability concerns: there is a fundamental friction between reusability of knowledge components representing sources of law and their relevance; Knowledge about prevailing and intended normative order is more relevant than knowledge of the structure and interface of legal institutions, but

1. the sources of law do not directly contain it,
2. it is less reusable in the sense of being very contextual, and
3. prevailing normative order changes due to changes in society that are not mirrored by changes in the sources of law.

7.4. Research Questions

The law intends to formalize a normative order, but normative order only exists if a large community of people generally make decisions in compliance with the normative order. Complying with the intended normative order, or recognizing non-compliance is however not trivial.

Central to this problem is often the problem of recognizing intentions (cf. section 4.1). Every action is intentional, directed towards causing something, but there is a distinction between intentional and unintentional consequences of one’s actions. The concept of intention was explained in terms of action: we recognize behaviour as action when we explain it as a *plan execution*. Since we recognize that cognitive resources are limited, we assume that people remain committed to plans over a longer time period. The decision to execute a plan takes place in a certain discrete time period

and consumes scarce cognitive resources. Every change or non-change that was part of the plan, regardless of whether one is the physical cause of the change, and regardless of whether the change was part of one's objective in planning, is in a sense intended.

I made a distinction between behaviour descriptions in law that do (ought-to-do) and do not (ought-to-be) explicitly address an agent role and action (cf. section 4.7). Even if no agent role is explicitly addressed, this does not change the fact that only action violates normative rules: the subject of a norm or a preference is in essence always a situated action under conscious control, even if the legislator had no specific script in mind or didn't bother to mention it explicitly, as in many technical rules for design and construction activities.

In order to come to the conclusion that someone made the wrong decision, the outside observer will if necessary decide 1) what plan was being executed, 2) when the decision to execute that plan was made, 3) which alternative plans were considered, and 4) generally, what the decision maker knows in a dispositional sense. This *folk psychology* model of decision making presents legal knowledge engineers with interesting, and important, research questions:

1. I have pointed out a fundamental *deadline* problem in relation to many obligations: it is not possible to determine whether an obligation to bring something about has been violated without setting a deadline, if one separates the obligation from an ascribed planning activity (cf. section 4.7.2). In this case and in the case of *negative causation*, also briefly mentioned, one violates a normative rule because of the alternative plan one *did* execute instead of those that comply with the obligation: it is not the case that one violates a rule by non-action. If this were not the case, setting deadlines would receive a lot more attention in the field of law than it does. Negative causation, despite its name, has very little to do with causation. Both passive and active behaviour can be considered to be action, if it is interpreted as the execution of a plan.
2. In relation to compliance dilemmas we also encountered the question of whether there is an alternative behaviour that is allowed. As pointed out in sections 6.3.1 and 6.7, it is not sufficient to believe that there must be such an alternative; One must also *know* one to be confident in the decision that some behaviour is disallowed. The alternative we choose is the best one we know, not the best one there is. This applies not only from an internal perspective, but also to the external observer who ascribes a choice between concrete alternatives to the decision maker. The compliance conflict in literature would be easy to formalize in terms of a discrete and finite set of alternatives in a menu. When disconnected from a theory of decision making that supplies such menus, it becomes a very hard knowledge representation problem.

The meaning of normative order is inseparable from a model of planning. While knowledge representation is usually interested in intelligent planning activity, in keeping with the philosophy of representing *knowledge as a productive factor*, the law is concerned with *all* planning activity, not only particularly intelligent planning. AI planning literature has little to offer: the norm may play a variety of roles in planning models. Depending on the purpose of planning, the norm may be an ordering on al-

ternatives, a reason to reject alternatives, but also very often the apparent reason why alternatives are not even imagined at all. The behaviour of others is usually predicted in terms of the prevailing normative order: non-compliance is not even foreseen, and certainly not planned for.

There has been some initial work on the task of *legal planning* (cf. for instance [265, 27]). These however also conceive of legal norms as mainly as constraints in planning. The role of norms as enablers of action is rarely considered.

Connection of planning and situation recognition to Schank's scripts and Kahnemann's reference transaction (cf. section 4.5) as a device for explaining a prevailing normative order is still an open research question, and an interesting one. The problem of planning and ascribing plans to others should certainly be a central problem in legal knowledge engineering.

The notion of normative order also presents us with other interesting research questions of a more practical nature. It for instance makes sense that certain organizational notions like *business process* can also be considered as conceptualizations of a prevailing or an intended normative order within an organization (depending on whether one is modeling an existing one or re-engineering one). A mapping between the language of law and normative order and the language of business process modeling would have practical commercial applications in organizational settings for the coordination of organizational change and the revision of deployed legal knowledge-based systems in reaction to changes in the sources of law.

7.4.1. Abstraction

The mechanisms of abstraction, aggregation, refinement, and ontological stratification discussed in section 2.3.1 also present us with interesting research problems, but explaining these forms of reasoning is not specifically relevant to legal reasoning.

In this book, these mechanisms have at some points been used as an explanation, but how they arise and how they work remains unexplained. Seemingly arbitrary choices of granularity level play a prominent role in two common sense reasoning subproblems in legal reasoning. These both involve relating legally relevant terms through common sense knowledge not supplied by the law.

The first one concerns the question of whether the behaviour or state of affairs sketched in one rule is more specific than the one in another. This issue is of interest to the defeasibility of indicative rules in general, and to the resolution of normative conflicts specifically. Specificity appears to be much broader than subsumption, and it is a problem for nonmonotonic reasoning in general (cf. section 6.3.1). Descriptions at completely different levels of *abstraction* are sometimes said to be in a specificity relation to each other: the solution to this question is perhaps found in the equally mysterious notions of aggregation and refinement as applied to abstractions, in the sense that refinements, which are more often than not *imperfect reductions*, are more specific than the thing they are a refinement of.

Trying to emulate this form of reasoning is a considerably less efficient investment of one's modeling effort than representing the results of validated arguments of this kind in court decisions and administrative guidelines as choice rules.

A second, similar issue exists for choice of granularity in the description of actions and events. As noted in section 3.5.3, whether something goes through a series of changes or remains the same over a time period, is mainly a matter of selective attention. One may choose different ontological strata, different abstractions at varying levels of granularity, when describing a behaviour, but grounding every description encountered in a fundamental causal model of reality is not a task the legal knowledge engineer can set himself when representing a source of law.

Causal reconstruction is however of interest in certain areas of law, mainly criminal law and tort law, and support for creating causal models may find a niche market in this context. There are interesting efforts in this direction in [152, 184], and indirectly in [23]. In most applications of law, the recognition of actions is not considered problematic and the question of causality does not arise. This is because we already recognized behaviour as action, i.e. ascribed a plan, and implicitly attributed responsibility for certain events – and non-events – to the planner.

Both issues concern “brute” *legally relevant* terms, and the more general problem of grounding these in a fully developed model of the relevant parts of consensus reality, sometimes included in the quest for *deep models* or *deep structure* of law as for instance found in [94, 20]. This notion of deep structure, and the benefits of uncovering it, has been implicitly adopted in this work.

The notion of the intended normative order in this book is a deep model to be uncovered. Some deep models are however maybe better left unaddressed. The problem of providing a “deep” explanation of specificity and causality is in the opinion of this author too hard and in practice too insignificant to consider them as requirements for practical legal reasoning. The explanation of the deep structure of law in [94] (also adopted by, amongst others, myself in [286]) does appear to assume that a fully developed model of the relevant parts of consensus reality, including the specificity question at least, is a feasible project, but our work in the last decade has not resulted in a satisfactory approach to uncovering it.

Reconstructing the limited set of reference transactions the legislator had in mind when he chose a certain legislative solution *must* be feasible. Specificity is only interesting if we find it in the form of an implied subsumption in the language chosen by the legislator, which only happens if the legislator *designed* the choice into his legislation. In this case it is naturally addressed by a choice rule (cf. section 6.6). Construction of fundamentally new conflicts should be left to the courts.