Joost Breuker and Rinke Hoekstra, 'DIRECT: Ontology-based Discovery of Responsibility and Causality in Legal Case Descriptions' in T. Gordon (ed.), *Legal Knowledge and Information Systems. Jurix 2004: The Seventeenth Annual Conference*. Amsterdam: IOS Press, 2004, pp. 59-68.

DIRECT: Ontology-based Discovery of Responsibility and Causality in Legal Case Descriptions

Joost Breuker and Rinke Hoekstra
Leibniz Center for Law, University of Amsterdam
{breuker,hoekstra}@lri.jur.uva.nl

Abstract. In this paper we present DIRECT, a system for automatic discovery of responsibility and causal relations in legal case descriptions based on LRI-Core, a core ontology that covers the main concepts that are common to all legal domains. These domains have a predominant common-sense character – the law is still for the people – and typical legal concepts such as norm, role responsibility, contract, etc. have a grounding in abstract common-sense conceptualizations. A common sense framework is even more necessary for the automatic discovery of the causal structure of legal case descriptions, since they often are even devoid of a specialized legal conceptualization. The paper presents a number of design principles that follow from the common-sense stance in developing the LRI-Core: the most important being cognitive plausibility. Furthermore, an approach is presented to enable the automatic analysis of cases described in terms of the ontology. Such analysis will determine causal chains in situation descriptions on the basis of which responsibility attribution can take place.

1 Introduction

Responsibility attribution is a vital element of reasoning about legal cases. As studied by Lehmann [9], (agent) causality, responsibility, liability and guilt are closely intertwined concepts in legal reasoning aimed at compensating for wrong doing. The work presented here is based on his work, but has a more formal approach. It aims at validating these and some new views by constructing a system DIRECT[1] that is capable of finding out which agents are to be held responsible for which actions, as described by some semantic representation of a (legal) case. A case description is an account of events and states (situations). Events are related, in particular by causal and intentional relationships. The test case we have in mind is an unrelated and unordered set of states, represented as instances of objects, processes, actions, agents, etc. DIRECT has to find out which of these states are causally related – and why – by inferring causal and intentional relationships (motives), i.e. DIRECT has to make 'sense' of the collections of events and states. We can compare the results with those of human reasoners. It should be noted that this project is not aimed at the construction of a kind of (semi-)automated judge. It is by far too early even to contemplate such a practical application. The aim of this study is to come to grips with the concepts that are central in legal reasoning (cf. e.g. [4]). A major question is which information, and how much, is needed in case descriptions for responsibility attribution. For instance, we want to test our strong hypothesis that causality is determined by classifying events as processes, and that temporal information about events is only a side-effect of time- and energy 'consumption' (See Figure 2). In fact, this study is to be viewed as an exercise in computational jurisprudence, rather than in legal knowledge engineering.

[1] DIRECT: DIscovery of REsponsibility and CausaliTy

Before we can actually perform these experiments, we have to prepare the knowledge bases to be used by DIRECT. The most important and general of these is a core ontology: LRI-Core. In conjunction with more specific extensions, it contains the definitions of the concepts that are needed to interpret (understand) the instances that make up the second knowledge base: the set of events and states that describe (occur in) a legal case. In this article we explain and justify the construction of these knowledge bases.

Cases are mostly expressed in the common-sense vocabulary people use in everyday communication. The legal norms, and definitions of legal concepts we commonly find in legislation are expressed using a mixture of these common-sense terms and (domain) specific legal jargon. We therefore need a common unifying vocabulary for reasoning on legal cases: a core ontology of law. That is what we will describe in summary in Section 2. In Section 3 both the philosophy of the experiments and the construction of the events knowledge base (case) is presented.

2 A Core Ontology of Law

A core ontology covers a 'field' like medicine, cultural goods, etc. which may consist of many (sub)domains. Law, for instance, is carved up in domains like criminal law, environmental law, private law, etc. A core ontology specifies the common conceptual denominators of a field, i.e. those abstract concepts that are part of all (or an important majority) of domains. In law concepts like 'norm', 'role', 'document', 'liability', etc. are present in all legal domains: implicitly or explicitly. A core ontology has several purposes. First, it enables re-use by allowing top down modeling of a new domain. This re-use role is very important in practical knowledge engineering as the core ontology becomes the unifying framework of multiple domain ontologies. Furthermore, a core ontology of law enables us to codify acquired expertise across the various domains we work on, i.e. it facilitates knowledge acquisition. Finally, as abstract core concepts play a pivotal role in reasoning, they may be a source for constructing special inference services, as e.g. in spatial and temporal reasoning. For instance, we developed a formalism and inference engine for reasoning with (legal) norms, as part of a legal core ontology for law [14].

Except for CYC[2], other foundational ontologies (e.g. SUMO[3], Sowa's upper ontology [12] and DOLCE [3]) do not take an explicit common-sense stance. Even in DOLCE, which is based upon human perception, cultural imprints and social conventions, the common-sense perspective is not explicitly developed. Furthermore, these other foundational ontologies do not (or hardly) contain conceptualizations of core concepts in law such as role, norm, responsibility, evidence. These considerations convinced us of the necessity of a common sense core ontology of law: LRI-Core.

2.1 LRI-Core

The common-sense stance in ontology development requires a different approach than that of other foundational ontologies. Where other approaches either work from premises in science (e.g. SUO) or language use (e.g. DOLCE, Sowa), the common-sense stance is founded on results from cognitive science. This means that the resulting ontology represents not exactly the structure of the world, or the terms we use in language, but rather how the world *appears* to exist to laymen.

In [1] we present arguments why we think that a common sense foundational ontology should focus around (at least) the five major categories of that form the top-layer of LRI-Core

[2] `http://www.cyc.com`

[3] Suggested Upper Merged Ontology; `http://ontology.teknowledge.com`

(see also Figure 1). The major line of argument is based on evolutionary principles. Primary conceptualizations are inspired by moving and sensing, i.e. real-life interactions with the *physical* world. The complexity of this causal world is reduced when we take a 'teleological' stance with respect to life, in particular on living organisms of the same species. A teleological or intentional stance implies that the actions of agents are assumed to be motivated by goals. Teleological reasoning works 'backward', i.e. it allows reasoning from end-states (goals) to current states. This is less complex than the branching of possible worlds in causal, forward reasoning. Living creatures seek the maintenance and reproduction of life.

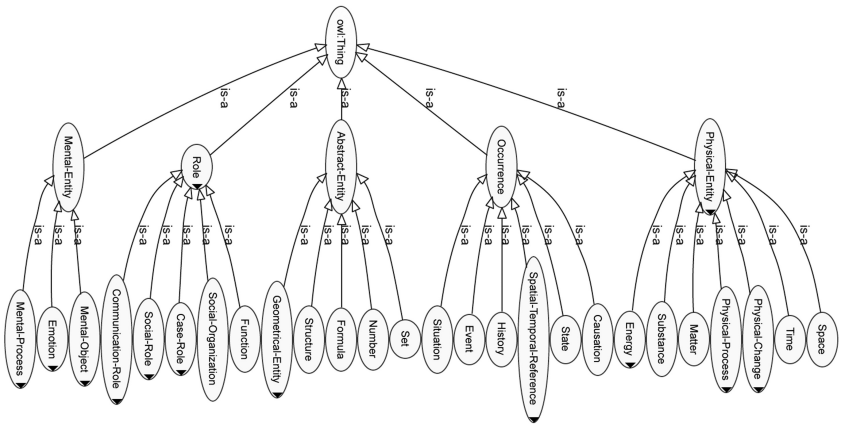

Figure 1: LRI-Core, top two layers

As human beings (and to some extent other higher mammals as well) discovered their own *mental* life, i.e. consciousness and self-awareness, the need arose for models of mental processes and objects. Awareness not only enables us to handle our own reasoning and emotions, but also to understand those of our fellow creatures in order to plan social activities and to communicate. Self-awareness enables 'reification', the building of metaphors that makes up *abstract* conceptualisations. These considerations convinced us that the mental world can be conceived as an intentional metaphor of the physical world, i.e. our mental life is made up of objects and processes. The categories we use to understand our own and other people's mental events mirror those of the physical world. The emergence of conscious planning and prediction of behavior has led to the conceptualization of *roles* that make up social organization. LRI-Core has thus been equipped with the following main categories: *physical*, *abstract* and *mental* concepts, *roles* and *occurrences*. Strictly speaking, occurrences are not part of an ontology, as we will explain below. LRI-Core is currently still under active development, it is expressed in OWL-DL, using the OWL-Plugin of Protg[4].

Occurrences An ontology should not be structured according to the *way* things occur in physical, mental, or fantasy worlds, but rather to *what* the things 'essentially' are. Ontology has a Platonic flavor in the sense that it specifies the ideas with which we understand a/the world as it passes by. Making sense of the world means that we build models of current, past, and even to some extent, future situations. The structure of entities occurring in a world is different from the (abstraction) structure of generic concepts that make up an ontology. The concepts defined in an ontology enable us to recognize entities and their relations as they occur in the world, i.e. they are the building blocks for the construction of actual situations and histories: partial models of real or imaginary worlds.

[4]See `http://protege.stanford.edu`

Histories describe the life line of individual entities, and situations are diachronic spatial structures of objects and processes. The distinction between situation models and the concepts we use to identify the elements (parts) of situations, is obscured in ontologies that make a fundamental distinction between occurrents (perdurants) and continuants (endurants) (e.g. Sowa, DOLCE). Perdurants are entities that have parts that change with time or place. "For example, the first movement of the (execution of) a symphony is a temporal part of it." ([3, p. 20]). In fact, if the authors say that the execution of a symphony is a perdurant, then all existences are perdurants. A stone, as a prototypical endurant in these ontologies, is in the execution of its life-line also perduring. Originally part of a rock, a stone may end up as sand on a beach, gradually spreading its parts spatially. However, when we take the concept of first-movement-of-a-symphony, or (pars-pro-toto) a symphony, there is nothing of temporary parts. Strictly speaking, all entities in situations are endurants; all concepts are perdurants.

The category of occurrences in LRI-Core captures those strictly temporal aspects related to the *execution* of scenarios involving objects and processes. This means that events and changes are occurrences, but processes are not. Where processes contain the information of the change(s) they bring about, events only describe a discrete difference between the situation before and after the event took place: they describe the input-output of the execution of a process, and are 'in' time. Furthermore, processes can make changes take place without there actually *being* a difference between before and after the process was executed. All this does not reduce the need for terms to talk about occurrences in general. For instance, above we have used terms like situation, event, history and entity. These terms refer to occurrences in an abstract sense that can legitimately be part of an ontology that defines concepts. Therefore, LRI-Core has a category of 'occurrences'.

Physical Entities The physical world evolves around two main classes: physical objects and processes. Objects are bits of matter, which in turn is typed by its substance. Objects have mass, extension, viz. form and aggregation state (limiting form). The existence of objects expresses the notion that matter (in particular solid matter) is what renders the physical world relatively stable and observable. Physical situations are usually described by the arrangement of instances of physical objects.

This intuition does not hold for processes. Processes consume energy to change objects, or parts of objects. Though naively problematic (See [5]), energy has conquered its place in common sense. Processes are described by the *changes* they bring about. Through interaction, processes can cause one another, leading to series of events that only stop at some equilibrium: in general conceived as that there are no interactions at all. In LRI-Core, processes are distinguished according to two views: (1) formal change (transformation, transduction and transfer) and (2) the kinds of (properties of) objects involved. (e.g. change of position, change of substance, etc.). A third property is whether a process produces or consumes energy.

The concept of process is often used as synonymous to action and activity. LRI-Core defines actions as processes that are initiated by an agent acting as actor. Notwithstanding the intricacies of mental (or agent) causation, the action itself is strictly physical: i.e. some muscle movement. The mental perspective implied by agent-causation is that actions are intended: they are preceded by some kind of intentional decision to act. Many ontologies use the term process to cover both processes and actions. Business processes, to give an example, consist of actions. By abstracting out the agents the work in an organization can be seen as consisting of anonymous processes. However, they do not exhibit the same causal gluing as in physical processes. For instance, business processes are planned and controlled, i.e. initiated by (supervising) agents. Analogous to planned activities, is the causal design of a device. Devices funnel causal chains of processes in such a way that they exhibit intended

behavior, i.e. their function.

Mental Entities Analogous to the physical world, we conceive the mind as consisting of (mental) objects, like concepts and memories, which are processed by mental processes that transform or transfer them. Memories are retrieved; concepts are formed. Moreover, these mental objects may be aggregations of more elementary objects. Memories consists of 'multi-media' representations of situations experienced; thoughts are made of more elementary parts like concepts.

The contents (substance) of these objects are representations. The conceptual content of thoughts is intended by propositional attitudes, like belief, desire, norm etc. Mental objects are processed or stored in containers (such as the mind) which in turn can have parts, e.g. memories. Mental processes like thinking, memorizing, imaging are operations on mental objects. The equivalent of physical energy in mental processing is the concept of emotion: the force that makes us focus our mental energies. There is however, an important difference between the mental and the physical. Where physical processes are governed by causation, mental processes are controlled by intention: i.e. they are *actions*. Thinking is seen as an action, since we assume that we have full control over our thoughts and can decide about what we are thinking. However, where our mind escapes our conscious intentions, e.g. forgetting an appointment, our stance is physical rather than an intentional.

The outcome of a mental process can be the intention to act, e.g. according to a structure of primary actions: a plan. These actions can be aimed at bringing about both physical and mental changes, e.g. changing the mental state of another agent. Such intended mentalistic actions are acts of communication (which also need some physical medium to transfer the intended mental state).

The role of mental conceptualizations is extremely important in understanding and communicating with other people. Their primary use lies in their role as building stones of models of the minds of other people: user-models. The intentional stance means that we attribute intentions and intention directed mental processing and belief to other people and to some extent animals (or even computers).

Roles Roles are entities in the mind, they do not 'really' exist. Roles are idealizations: we may not play a role correctly. An important distinction should be made between playing a role and the role itself: "agents can act, and roles cannot" [11]. Correcting incorrect role playing does not mean that we change the role: we change our behaviour. Like plans and processes, roles in ontologies are often confounded with their execution, in the same way as the execution of a symphony may be confounded with the symphony itself. The original meaning of the term role refers to a role of paper that contained the text of an actor in a play. Also the role-taker (some agent) and the role are often confounded, which may become obvious when we identify a role with a person. These kinds of confusions have made conceptual modelers aware of the tricky issues about roles (see e.g.[13]).

Roles are often viewed as relationships ([12, 13, 10]). Indeed, social roles have mutuality and complementarity. No students without teachers; no parents without children; no speakers without hearers, etc. In theory of law, a related view exists about the mutuality of legal positions: i.e. rights and duties [6, 7]. For instance, if citizens have the obligation to vote, the government has the duty to enable this voting. Nevertheless, this complementarity of roles might not be of enough importance to grant their representation as relationships in an ontology. The ontology may specify such relationships, but the primary notion of role is as a concept.

This becomes clear when we look at roles as concepts, i.e. at what roles *mean*. Roles are behavioral requirements on role execution and on qualifications of role taking. These

requirements are prescriptions, i.e. they are normative. In modern society many roles have formal requirements enforced by *law*. Legislation addresses actors by the roles they play[5]. If actual behaviour deviates from the norms attached to these roles we violate the law. Violations are based upon the distinction between the prescription (role) and role performance. Therefore, in court, it is the actor of the role who is made responsible: as a person; not as a role. Even the fictitious concept of legal-person for social organizations turns into concrete responsibilities of the liable persons who have mis-performed their roles.

Abstract Entities As all concepts are abstractions, one may argue that a separate abstract world is difficult to see. However, common sense knows about a (small) number of proto-mathematical concepts, such as collections, sequences and count-numbers (positive integers). [8] even argue that these common sense notions are the real roots of our mathematics. Nonetheless, this kind of semi-formal abstractions do not play a very central role in law, and therefore LRI-Core is thinly populated with abstract classes.

2.2 Uses of LRI-Core

The ontology is developed as a true 'core' ontology, aimed at being the unifying semantic backbone of multiple applications: a system for predicting behavior of citizens to new or changed regulations; a system for automatic detection of normative conflicts; a conceptual retrieval system for a large (legal) document-base, a system for legislative drafting, a system for normative assessment of legal cases[6]. Just as the development of some of these systems indicated the necessity of a core ontology of law, the development of future systems forms its ideal testing ground.

3 Ontology-based Causal Discovery

For a full understanding and representation of legal reasoning we need an explicit representation of what a case is *about*. This requires not only the representation of the subject of a case (its contents), but also the development of a formal framework for describing (legal) situations, causal relations, responsibility relations and normative constraints on situations. DIRECT represents our conviction that the causal reconstruction of a case is a necessary (though insufficient) requirement for the attribution of responsibility. A purely legal stance on liability is not sufficient for explaining situations in legal cases. Whether someone is liable in a certain case, is for a large part defined in law. However, a causal account is minimally necessary for discriminating between cases in which a strict legal (causal minimalist) perspective is sufficient and cases in which it isn't. Furthermore, we argue that it is more practical to use a causal representation of a case than a preliminary legal qualification for discriminating between these cases. We will come back to this issue in more detail in Section 3.3.

3.1 A Causal Reconstruction

Before a case is presented to the legal system it lacks the legal qualifications (or even legal terminology) which are common in the prototype example of a 'legal case'. An important purpose of DIRECT is to establish the amount of information input that is *minimally* needed

[5] An exception to this rule is in criminal law.

[6] For more information regarding these applications we refer to our website at `http://www.lri.jur.uva.nl`.

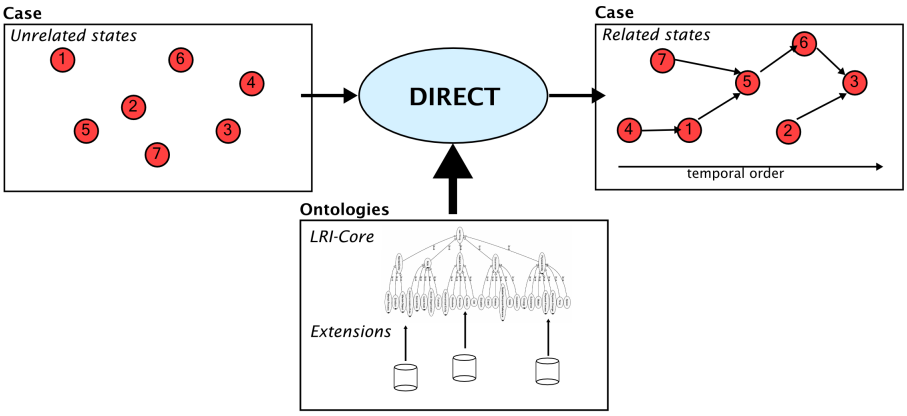

Figure 2: Causal reconstruction of a case in DIRECT

for legal reasoning on situations. Figure 2 shows how a minimally represented case is reconstructed by DIRECT. Cases are presented to the system in their most rudimentary form: an unsorted collection of situations, i.e. the situations are not ordered according to their occurrence in time. The system tries to find an appropriate ordering and identifies causal relations between the events on the basis of information stored in an ontology base. A description of a situation (or state) is a set of individuals, their attributes and relations. Individuals are instantiations of objects (in the broad sense) defined in LRI-Core or one of its extensions, e.g. the concept 'knife' is not central to the domain of law and is therefore not part of LRI-Core. Cases can be explained according to two perspectives: a strictly *causal* perspective and an *intentional* perspective.

3.1.1 Events and Processes

The changes we find between situations are events, i.e. events are identified by the corresponding differences between situations. Temporal ordering is reconstructed through the principle of *minimal change*: situations that differ in very few aspects are more likely to be close in time, than situations with more differences between them. DIRECT works analogous to video decompression: a cut-up film (case description) is resynched by reconstructing the minimal changes possible for each frame (situation). The more complete situation descriptions are, the easier it becomes to identify the events between the situations.

Changes, that is events, can only occur when some kind of process takes place. As processes are defined by the changes they bring about, they also provide the (causal) explanation of why certain events take place. Recognizing causal relations is thus largely about recognizing processes and actions, i.e. it is about *classifying* the state-changes in a case description as processes or actions. Processes can of course span multiple state-changes. Furthermore, the location of objects in situations plays an especially important role as location limits causal propagation: events are bound by location.

3.1.2 Actions and Intentionality

The previous section we have concerned ourselves primarily with *physical causation* (or causation in fact). However, this form of causation is relatively straightforward compared to its intentional counterpart: *agent causation*. The recognition of agent causation is an important prerequisite for responsibility attribution. Agent causation is about recognizing events which are brought about by actions of intentional agents. Two variants of action can be discerned: *physical* actions (strict agent causation) and *communication* actions (interpersonal

causation) (cf.[4]). Physical actions are operations by agents on the physical world, i.e. it concerns changes caused by (intentional) muscle movement. Where it differs from physical processes, is that these operations are guided by some instigating motive. The agent has some intention of performing that action in order to achieve some goal. Communication actions, on the other hand, do not have any significant physical influence, but they rather influence the mental state of other agents, i.e. they are purely interpersonal. Communication transfers intentions (motives, plans) and beliefs to other people. Some agent can provide reasons or draw attention to reasons which influence the conduct of another agent, who causes some event ([4]). Interpersonal causation is thus tightly connected to the notions of coercion and authority.

For the detection of agent causation, DIRECT thus needs to make a number of assumptions about the mental models of agents participating in the case. It has to keep track of events left implicit in the situation description to be able to infer the intentions of agents. For this it uses a default mental model to keep track of the things agents might have seen, heard or otherwise internalized.

3.1.3 Equilibria, Attempts and Negligence

The ontology provides a typology of processes and actions, i.e. a typology of templates to which events or differences between states are matched. However, some elements of legal cases are inherently 'hidden'. Some phenomena do inhibit change, but do not create differences, i.e. no events take place. Chemical equilibria, for instance, are important in explaining certain occurrences but do not themselves constitute any discrete change in the state of an object. *Attempts* to perform an action do not result in the desired action, but rather some rough approximation of it. These 'partial' actions can be explained by the intention (or plan) to perform the action, and a match of the actions that were actually performed with part of the event-structure of the intended action.

The concept of negligence, sometimes titled *negative causation* ([9, 4]), cannot be covered without taking legal considerations into account. When someone neglects to act where the possibility to act presented itself, this can only be deemed some form of causation when the outcome of the event-structure was designated to be undesirable. This moral (or legal) qualification is not part of the neutral description of a case, and as such is part of the attribution of responsibility.

3.2 Inner Workings of DIRECT

Figure 2 shows the causal reconstruction process performed by DIRECT. As mentioned in Section 3.1, the input consists of an unsorted set of situation descriptions. Situation descriptions are essentially conjunctions of states. These states are indexed by the name of the entity they refer to. Every entity (agents and objects) have a unique identity across situation descriptions. The system uses this unique identity to build histories of every entity. For each possible state transition, the system checks for each entity the differences between the two states. DIRECT identifies the events characterized by the differences (changes) between these situations. Processes, characterised by the events they bring about, provide the causal structure along which a case scenario unfolds itself. Actions are those processes which are initiated by agents. Intentions play a vital part in the identification of actions undertaken by the participants in a case.

Simply put, this works by matching differences to a typology of processes and actions, stored in an extension to LRI-Core. If no single matching action or process can be identified, the system assumes an underspecification in the state description and tries combinations of

processes[7]. The typology of processes and actions needs to be very elaborate to be able to capture the complexity of dynamics in the world. Furthermore, the representation of complex intentional processes but also knowledge about the beliefs of others, knowledge of situations, planning etc. will put a fierce strain on both the representation formalism (OWL-DL) and the inference engine used.

3.3 Responsibility Attribution

This paper presents an approach to responsibility attribution in which a causal account of a situation is taken as input for the attributive process. The definition of [4] states that legal responsibility is the liability of a person to be punished, forced to compensate, or otherwise subjected to a sanction by the law. In other words, if a judge or other authorised party acknowledges someones liability to be punished etc. the person is considered to be held legally responsible. In that light, our ambitions do not reach all the way: we intend to determine whether someone is *potentially legally responsible.* Grounds for the attribution of legal responsibility to a person for a given harm are: the conduct of a person, the causal connection between the conduct of the person and the given harm, the fault legally implied by the conduct of the person ([4]).

According to causal minimalists, liability is attributed strictly on the basis of comparing the agents participating in a situation with a set of norms. What really 'happens' does not play any role of significance. However, as we have argued in Section 3, very often cases require a (common sense) causal explanation for proper attribution. Besides such a (physical) causal explanation, intentional explanations (agent causation) are often necessary to be able to distinguish between accidents and purposeful actions.

A minimalist approach would require a system to test for every case whether exceptions to the default rules hold. We reject this approach because it presupposes that the law covers every possible exception. It does (in our view) not correspond to the way people work, and it requires the case description to be in legal terms. To our conviction minimal requirements for responsibility attribution are: 1) the identification of causal (and intentional) relations in a case description, and 2) the encoding of general principles of liability attribution in law. Once the causal structure of a case is identified, legal responsibility attribution comes down to establishing whether a legal norm has been violated, and following the causal chain backward until we encounter the primary 'causer' of the harm. Principles of liability attribution can then be applied to determine whether this primary causer is to be held responsible for the norm violation. Thus, in order for a defendant to be held responsible for a crime or tort, it must be proved that a sufficient causal link relates the defendant's actions to the criminal event or damage in question.

4 Discussion and Future Work

We presented a system for ontology-based discovery of causal relations in legal cases for the purpose of responsibility attribution. This system will be used to study the concepts used in this process, and the minimal requirements for reasoning about legal cases. The system depends on an ontology, LRI-Core, and some extensions for the terminology and structures used to represent cases. LRI-Core is a common-sense core ontology of law. This common sense stance distinguishes LRI-Core from other foundational ontologies. Based on psychological evidence, it is closer to our everyday interpretation of the world around us than other more science, or philosophy-based approaches. It is our conviction that representing

[7] During experimentation, a possible interface to DIRECT could allow the user to repair state descriptions or the ontologies by entering new knowledge.

knowledge in and on the legal system requires such psychologically plausible common sense conceptualization of the world.

As we speak, DIRECT's development is still very much in the early stages. Most of the work described in this article is the result of combining and operationalizing the approaches developed earlier in [1, 2, 9, 14]. Future work includes the development of an extension to LRI-Core which contains a typology of processes and actions. This development will be a two-step process. We will start experimenting with very simple relations in block-world like scenarios. Once the general principle is tested and running, we will extend the matching system and typology to cover more complex scenarios which include elaborate intentions such as plans, coercion, persuasion etc.

References

[1] Joost Breuker and Rinke Hoekstra. Core concepts of law: taking common-sense seriously. In *Proceedings of Formal Ontologies in Information Systems (FOIS-2004)*. IOS-Press, 2004.

[2] Joost Breuker and Rinke Hoekstra. Epistemology and ontology in core ontologies: FOLaw and LRI-Core, two core ontologies for law. In *Proceedings of EKAW Workshop on Core ontologies*. CEUR, 2004.

[3] A. Gangemi, N. Guarino, C. Masolo, A. Oltramari, and L. Schneider. Sweetening ontologies with DOLCE. In A. Gomez-Perez and V.R. Benjamins, editors, *Proceedings of the EKAW-2002*, pages 166–181. Springer, 2002.

[4] H. Hart and T. Honoré. *Causation in the Law*. Oxford University Press, New York, second edition edition, 1985.

[5] P. J. Hayes. The second naive physics manifesto. In J. R. Hobbs and R. C. Moore, editors, *Formal Theories of the Common Sense World*, pages 1–36. Ablex Publishing Corporation, Norwood, 1985.

[6] W. Hohfeld. *Fundamental Legal Conceptions as Applied in Legal Reasoning*. Yale University Press, 1919. Edited by W.W. Cook, fourth printing, 1966.

[7] Hans Kelsen. *General Theory of Norms*. Clarendon Press, Oxford, 1991.

[8] George Lakoff and Rafael Núñez. *Where Mathematics Comes From*. Basic Books, 2000.

[9] J. Lehmann. *Causation in Artificial Intelligence and Law: a modelling approach*. PhD thesis, University of Amsterdam, 2003.

[10] C. Masolo, L. Vieu, E. Bottazzi, C Catenacci, R. Ferrario, Aldo Gangemi, and Nicola Guarino. Social roles and their descriptions. In *Proceedings of the Knowledge Representation Workshop*. 2004.

[11] O. Pacheco and J. Carmo. A role based model for the normative specification of organized collective agency and agents interaction. *Journal of Autonomous Agents and Multi-Agent Systems*, 6:145–184, 2003.

[12] John F. Sowa. *Knowledge Representation: Logical Philosophical, and Computational Foundations*. Brooks Cole Publishing Co, Pacific Grove, CA, 2000.

[13] Friedrich Steimann. On the representation of roles in object-oriented and conceptual modelling. *Data and Knowledge Engineering*, 35(1):83–106, October 2000.

[14] A. Valente. *Legal knowledge engineering: A modelling approach*. IOS Press, Amsterdam, The Netherlands, 1995.