

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

On the ontological status of plans and norms

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/99629> since 2016-11-30T23:27:04Z

Published version:

DOI:10.1007/s10506-005-4158-x

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

On the ontological status of plans and norms

GUIDO BOELLA, LEONARDO LESMO and ROSSANA DAMIANO

Dipartimento di Informatica, Centro di Scienze Cognitive, Centro per l'Ontologia Teorica e

Applicata, Università di Torino, Torino, Italy

E-mail: guido@di.unito.it

Abstract. This article describes an ontological model of norms. The basic assumption is that a substantial part of a legal system is grounded on the concept of *agency*. Since a legal system aims at regulating a society, then its goal can be achieved only by affecting the behaviour of the members of the society. We assume that a society is made up of *agents* (which can be individuals, institutions, software programs, etc.), that agents have beliefs, goals and preferences, and that they commit to intentions in order to choose a line of behaviour. The role of norms, within a legal system, is to specify how and when the chosen behaviour agrees with the basic principles of the legal system. In this article, we show how a model based on plans can be the basis for the ontological representation of norms, which are expressed as constraints on the possible plans an agent may choose to guide its behaviour. Moreover, the paper describes how the proposed model can be linked to the upper level of a philosophically well-founded ontology (DOLCE); in this way, the model is set in a wider perspective, which opens the way to further developments.

Key words: legal knowledge representation, ontologies, planning, BDI agents

1. Introduction

The role of ontologies in computer science is to provide a formal description of the world. They aim at enabling a machine to find out relevant information from documents, to recognize described situations, to draw conclusions on the basis of known facts, etc. In this context, it is not relevant that the description refers to the world 'as it is' (true ontology, in the classical view), but what is important is that it represents the world as it is conceived by sentient beings (epistemology): so far as an ontological structure is embodied in a computational agent, it represents the way that agent sees the world, but to the extent where that structure is shared among a community of agents, it represents the way the community sees the world, and if the ontology is also used in man/machine communication, it must share with humans some basic ideas about the way the world is organized (Breuker and Winkels In press).

So, an ontology must be properly founded, in order to make possible this multiple sharing. But often ontologies are used in applications concerning

specific domains, so that the proper foundations must be linked to domain conceptualizations. However, it is not clear how this link can be achieved. A first problem concerns the layman/expert point of view. Obviously, a physician has a very different view of what is a human body (and how it works) than a non-physician. So, it may be assumed that an ontology devoted to the medical field must be able to express the view of the physician, in order to put at her/his disposal, when needed, the vast body of knowledge concerned, for instance, with the liver, or with what an illness is.¹ On the other hand, a physician must be able to talk with her/his patients: this often requires a rephrasing of complex medical descriptions in (possibly imprecise, but understandable) non-technical terms. And s/he must also be able to understand what are the symptoms, even if they are not described by the patient in the most proper way. This happens in many fields, and it is especially important in the legal field, where many legal concepts are inextricably linked to the commonsense view of the world. This produces two, partially related, effects: first, in legal documents, as well as in the Statutes, there may appear terms (as window, water, tree), which are not defined in the legal systems, so that it is assumed that the definition is borrowed from a general ontology; second, some concepts, which are of common use (such as property, person, duty), are changed by a legal system by means of a body of definitions and norms appearing in the system, so that the jurist's view of a 'person' is not exactly the same as the layman's view.

The goal of this paper is twofold: first, to show how commonsense concepts (in particular, action and behaviour) can be linked to legal concepts (in particular, norms and rules); second, to show how the resulting structure can be linked to a properly founded top-level ontology. It must be clear that we do not aim at covering the ontological status of all types of rules. For instance, in a legal system, there are rules concerning the definition of concepts (e.g. Non Governative Organization) and rules concerning the way a given institutional role can be established (e.g. the Secretary of an NGO), and which are its attributions. We are not concerned with definitions, since we assume they are given in advance, and already included in the ontology. But we are (at least partially) concerned with institutional roles and their attributions, since institutional roles, as authorities, are a (predefined) category of people, and their attributions define what they can and cannot do. Also, we are not concerned with the way a Court gives a verdict, which is a very complex process, but we are concerned with the presence, in the ontology, of all the concepts required to describe the facts and the evidence available to the court. In other words, this paper faces a small but relevant fraction of the legal knowledge:² it covers the norms of conduct, i.e., the norms that specify what a legal actor can and cannot do. The claim is that these rules express constraints on actions, and their relevance stands on the view that a legal system aims at organizing a society in such a way that the rights and duties of

individuals and organizations are properly balanced, on the basis of superior principles. And these rights and duties are expressed in terms of norms of conduct. The article shows that the same approach can be applied to the modelling of legal powers; in this case, the constraints refer to the plans that must be executed to establish new rights and duties.

We must point out that we do not aim just at describing norm-related concepts, but also at showing how norms affect behaviour. In other words, the formalization of prescriptive rules is assumed to be available to an agent that must do things, and that must choose what to do: this agent uses the rules to decide what he will do next; in principle, he will choose a line of action that complies with the existing rules, but in some cases he may decide that a rule is not worth being respected, either because it is in conflict with other norms, or because the risk of being sanctioned is low. Although this behaviour is only partially moral, it is the way many rational agents choose to act: it must be properly modelled, since it explicitly includes the concept of sanction, and, implicitly, guides the legislator to properly choose the sanction associated with a given rule.

The paper is organized as follows: the next section presents the upper-level of the ontology (DOLCE), and our formalization of plans; in the third section, the representation of norms is described; the fourth section shows how power and authorities fit our approach; in Section 5, we present a case study addressing the legal concept of ‘fruit’; then, we compare our approach with other ontological models of the legal systems and, in Section 7, we present some ontological models of plans; the conclusions close the paper.

2. The upper-level of the ontology, and the plans

In order to introduce our model of plans and norms into a well-founded ontology, we have chosen, as our starting point, the DOLCE ontology (Gangemi et al. 2002). The main advantage of DOLCE is the clear philosophical foundation of its basic categories: they have been selected as the result of a research effort that has lasted for some years, and the motivations for the final choices can be found in the referenced papers, and can be traced back to various previous proposals (Guarino 1998; Gangemi et al. 2001; Welty and Guarino 2001; Guarino and Welty 2002). Most of the structures addressed in the paper are based on the KIF implementation reported in (Masolo et al. 2003).

2.1. DOLCE

Without entering in any detail (the interested reader is addressed to the original papers), we report here some brief comments on the top-level of

DOLCE, which seem necessary in order to explain the connections to plans and norms. Some of the concepts defined in DOLCE are depicted in Figure 1.

The upper-level categories are *Abstract*, *Quality*, *Situation*, *Endurant*, and *Perdurant*. The assumption about the existence of abstract entities and qualities is rather common (with the known philosophical intricacies). The distinction between *Endurants*³ and *Perdurants* can be specified by citing the authors:

“*Endurants are wholly present (i.e., all their proper parts are present) at any time they are present. Perdurants, on the other hand, just extend in time by accumulating different temporal parts, so that, at any time they are present, they are only partially present, in the sense that some of their proper temporal parts (e.g. their previous or future phases) may be not present*” (Gangemi et al. 2002, Section 2.1).

In our words, which could not be shared by the authors, this means that the most relevant structuring of endurants is in space (so they extend in full across time), while the most relevant structuring of perdurants is in time (so they are cut in parts by the flow of time).

Situation refers to states of affairs (SOA). According to the authors, other names used in the literature for a *Situation* are setting, Gestalt, or configuration. Example of *Situation* are a clinical condition, a murder, and a speed excess case. It is interesting to note that in the Documentation of *Situation* [in the KIF knowledge base (Gangemi et al. 2003, p. 195)], it is stated that situations “*could be equivalently modelled as a special complex perdurant defined through its relations to qualities, regions, and endurants*”. *Situations* satisfy *Descriptions*, in particular *S-Descriptions* (situation description), see

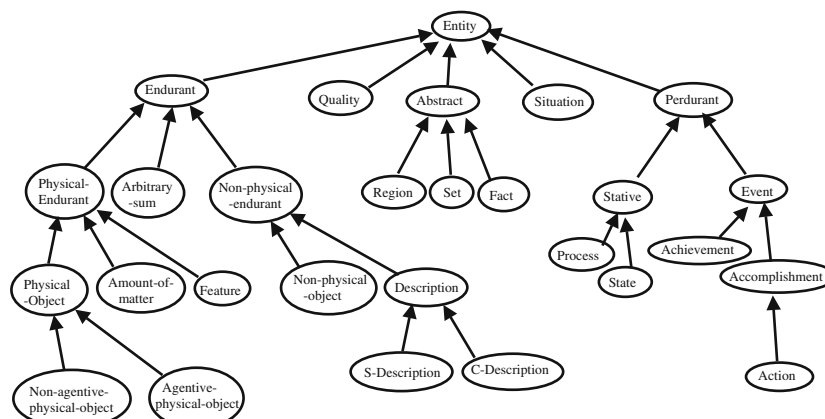


Figure 1. A portion of the upper level of DOLCE [from Masolo et al. (2003)].

Figure 1. In our interpretation, a situation is the ‘actual’ counterpart of a representation (description), i.e., a piece of reality, if one is a realist.

2.2. PLANS

Descriptions, which are non-physical endurants, are particularly important in this article, since, in DOLCE, plans are *S-Descriptions* (see Figure 2). So, *Plans* are particular descriptions of situations. This placement reflects a view where plans describe something; although we would rather take a plan as the situation which is described, in order to maintain the compatibility with DOLCE, we focus in this article on the description view. However, we must point out that we do not consider an actual plan as something occurring in the world, but rather as a ‘mental’ situation. It consists in the abstract object (the ‘idea’) that is associated with ‘potential’ occurrences of events. Consequently, DOLCE’s *Plan* is taken as the mental description of this object, a description which (in DOLCE’s terms) is *represented-by* a *MentalInformationObject* which, in its turn, is *realized-by* a given configuration of brain neurons. This is a very strong ontological commitment, which would require a lengthy discussion which, for the sake of brevity, we must omit here.

What is known about *Plans* in DOLCE is that they are composed of *Tasks* and include some *Functional-Roles*; moreover, there is a *Planning-activity*, which is *sequenced-by* various types of *Tasks*. Actually, we must be clear about what we mean by *Plan*. In particular, we must keep apart three different views on plans:

1. The description(s) which an agent is reasoning about while he is trying to decide what to do (one or more partially specified *plan schemas*).
2. The *intention*, i.e., a single, possibly still underspecified, sequence of steps which the agent has decided will drive her/his behaviour.
3. The *execution of a plan*, i.e., the actual, real-world, events that the agent carries out.




Figure 2. The placement of *Plan* in the ontology.

What appears in Figure 2 is the first view of plans, so that it can be imagined as a tree-like structure, where some decisions have been already taken, while others are suspended. For example, I may know about plans for travelling to far cities. And I may have the goal of being in Rome tomorrow. So I can examine the plans “go to a far city by train” and “go to a far city by plane” (two different plan schemas). If I have chosen the second schema, then I have certainly fixed the ‘airport of arrival’ and the ‘airport of departure’ (i.e., two of the participants in the plan), but, possibly, the departure time is just given as ‘some time in the morning of tomorrow before 8.00 a.m.’, i.e., an underspecified description of the ‘departure time’ participant. Also note that a different plan is obtained when I have reserved the seat, since at that time the “departure time” has been fixed: this new plan is a plan ‘more specific’ than (*refines*, in DOLCE’s terms) the original one (see Figure 3).

With respect to Figure 2, we note that:

- Apart from the standard *subcategory* arcs (thick arrows), some other arcs have been introduced. They refer to *relations* in the definition, and are shown as thin arrows with a small box. The shapes that substitute the boxes in *p-preconditions* and *p-effects* are just a shorthand, to keep more compact the next figures. Each *relation* has a name shown near the box; the arrow enters a category, which specifies the range of the *relation*; this is an adaptation of the KL-One notation (Brachman and Schmolze 1985), which complies with the KIF implementation of DOLCE; in Appendix A, we report an example of how a standard ontological language (i.e., LOOM MacGregor 1991; MacGregor and Brill 1991) can be used to express the concepts and relations appearing in the drawings.




Figure 3. Plans and more-specific plans.

- The range of *p-preconditions* and *p-effects* is *S-Description*. This is a partial description of a state; it can be seen as an underdetermined state, where some characteristics are fixed, while others are left unspecified. In DOLCE, *precondition* is a relation between *S-Descriptions* and *situations*. We believe that is more homogeneous to maintain a relation between *S-Descriptions*; of course, the *S-Description* which acts as *p-precondition* will be *satisfied-by* DOLCE's *precondition*. For example, the plan of going Rome by plane is executable if the *S-Description* "The Rome Airport is open" is satisfied by the current *Situation*.
- *s-participant* specifies the entities involved in the plan; its range is *PartiallyDescribedEntity*. This is the role taken by a *C-Descriptions* (i.e., the *Descriptions* which describe anything which is not a *Situation*) in a planning context; *C-Descriptions* describe entities having some properties, or being in some relations with other known entities (the event that happened to me in some place at a given time, the platypus I saw in the Sydney park). Participants in plans need not be completely specified (I will take *a taxi*), but they can be described more or less precisely (I will take *a yellow taxi with an old, experienced driver*). The extension of a *PartiallyDescribedEntity* is a set of *Endurants*, such that the actual participant in the plan must belong to this set. The dotted oval associated with *PartiallyDescribedEntity* refers to its being a *material role* (see Guarino and Welty 2000): no new entity has come into being, but an existing one plays a new role (as when a person enrolls in a University s/he takes the material role of student).

Of course, many features of the plan ontology have been left out, but the goal of this paper is mainly to address the role of norms in the planning framework, so only the concepts necessary to enable us to specify the relationship between plans and norms have been described.

2.3. PLANNING

As stated above, plans are the objects that are managed in the process of planning. In Appendix B, the interested reader may find a sketch of an ontological representation of how planning may be carried out. In Figure 4, instead, we refrain from going into the structure of the planning action, but we simply represent the elements that affect that action. Intuitively, the idea is that *Thinking* is a *Perdurant* which has as participants the descriptions embodied in the brain. In particular, *Planning* is that type of *Thinking* that works upon the mental representation of plans. It produces new *Plans* by considering⁴ known *Plans* (going to Rome by bus), by specializing existing *Plans* (the taxi company has been chosen), or by discarding *Plans* (see the comments about the PLANET system in Section 7).




Figure 4. The Planning activity.

What is shown in Figure 4 is just the final outcome of this process, i.e., the fact that a plan has become an *Intention*. No *Intention* concept is currently defined in DOLCE, but *has-BDI-on* is a *modality-target* relation holding between an *Agent-role* and a *course* (of events). Since *Plans* are composed of *Tasks*, i.e., a specific type of *Course*, and who plans plays an *Agent-role* with respect to the planning action, we call *Intention* the role assumed by the sequence of *Tasks* which the planning agent has chosen. Correspondingly, we call *has-intention-on* the subrelation of *has-BDI-on* which refers to intentions (wrt. beliefs and desires). Although this is the actual picture in DOLCE (and we comply with it in Figure 4), in order to simplify the subsequent figures, we will merge the *Plan* and the associated *Complex-Task*. Consequently, in the following, we assume that *Plan* embodies also the *Tasks* it is composed of.

Notes to Figure 4:

- There are two effects of *Planning*: a *Complex-task* associated with a *Plan* has become an *Intention*, and the planning agent holds that intention (*has-intention-on*). The black rectangle on the *effects* arc is just used to keep together the various effects (it is just a logical And); the “+” on the arcs exiting the black rectangle mean that these are *positive* effects: they did not hold before the action, and they hold after the action.
- The *in-mind-of* relation holds among a (mental) *C-description* and an *Agent-role*. It means that the *Endurant* playing that role has in its brain a physical representation of that description.
- *Intention* is a *material role* that is played by the *ComplexTask*. The circle on the arc has no special function, except as a graphical means to specify that that arc is one of the effects.

- The diamond on the arc exiting from *Intention* is a constraint on the relations: it specifies that the two elements identified by the dotted arrows exiting the '=' box must be the same individual. The first of them is the *Agent-role* such that the entity playing that role 'has in mind' the *Intention*, and the second is the *Agent-role* of the *Planning* which has produced that *Intention*.⁵ In other word, in this model, agents plan for themselves (*p-effect-of* is just the inverse relation of *p-effects*).

What is not represented in Figure 4 is the role of *Intention* in driving behaviour: agents always act; so, if the current time is t_i , then agent *A* is doing something at t_i and will try to do something at t_{i+1} ; what is that something is specified by *A*'s current intention.

2.4. UTILITIES

Since, during planning, an agent considers many *Plans*, but just one of them becomes his *Intention*, he needs some way to decide which one of two (or more) *Plans* is the best. The model presented here is based on the assumption that:

- The choice is based on the agent's *preferences*;
- The preferences are applied to states that can be originated (either directly or indirectly) by the execution of the plan;
- From a computational point of view, applying the preferences to a state amounts to evaluating the *utility* of that state;

In our view, these three criteria characterize the *rationality* of agents. The use of utilities in the planning activity has been described elsewhere (Boella and Lesmo 2002). Here, some comments are needed on the second point. In fact, the reason why the preferences are not necessarily applied to (i.e., the utility is not necessarily evaluated on) the state directly resulting from the execution of the plan, is one of the key features of our approach to norms.

In general, the execution of a plan has certain effects on the environment (as specified by the plan *p-effects*). But the new situation that obtains can also produce further foreseeable changes. This is especially true in case the world is populated by other agents, who may interact with the planning agent. So, it seems reasonable to try to imagine how the world would be after some sequence of reactions from the environment. This is similar to what happens (in a very simplified way) in two-players games, where a player must evaluate his moves in terms of the subsequent possible moves of the opponent; it is clear that the ability of foreseeing in advance the evolutions of the game (for two or more moves) is one of the keys for finding a successful strategy.

So, the utility of a plan is evaluated on the basis of the state that obtains after the execution of the plan, and after some re-actions of the environment to the state that obtains (we assume here a two-level lookahead: plan

execution + reaction). The major point is that the environment is assumed to be able to react to the breach of norms; so, in order to model the behaviour of an agent, what is needed is a way to specify how the environment reacts to such a situation. Note that we have used the very general term ‘environment’; this leaves unspecified which are the relevant component of such an environment. Currently, we assume the existence of a *Normative-Agent-Role*, who has to enforce the respect of law. However, this agent can be a single individual (a policeman) or a social body (the set of other agents surrounding the planning agent). Moreover, the presence of normative agents does not exclude that the respect of a norm is a value ‘per se’, i.e., it is embodied in preferences of the agent.

3. Norms

Since law is a social construction, its bases lie on the interaction among agents. It has convincingly been argued that these interactions can be governed by some ‘legal relations’, and that all legal relations can be expressed in terms of four basic relations (Hohfeld 1919; Allen and Saxon 1995). So, we start our presentation by showing how the basic legal relations fit our overall framework (see also Boella et al. 2001). The four legal relations are: DUTY, RIGHT, NO-RIGHT, and PRIVILEGE. They are defined in terms of two interacting agents and of an action, in the following way:

1. X has the RIGHT that Y performs A, iff in case Y does not perform A, X may react with an action R such that Y is negatively affected by R.
2. Y has the DUTY to perform A for X, iff the same condition as the one above holds (i.e., DUTY is the “opposite” of RIGHT, where the roles of the two agents are interchanged).
3. X has NO-RIGHT that Y performs A, iff in case Y does not perform A, X may not react with an action R such that Y is negatively affected by R (so NO-RIGHT is the “inverse” of RIGHT, where the possible reaction is not allowed, i.e., negated).
4. Y has the PRIVILEGE to perform A with respect to X, iff in case Y does perform A, X may not react with an action R such that Y is negatively affected by R (so, PRIVILEGE is the “converse” of the “inverse” of the “opposite” of RIGHT, where we take “converse” to refer to the execution of A by Y, instead to the failure to execute A, as in all other relations).

Three comments. First, there are three parameters that come into play: the perspective (i.e., if we look at the relation from the point of view of X or Y;

what Hohfeld calls “opposites”), the possibility to react (“inverse”), and the fact that the definition refers to Y doing or not A (what we have called “converse”). It seems that “converse” is not strictly required, and it is introduced just to conform to a linguistic view (the existence of the term “privilege”). And in fact, just four relations are defined and not all eight combinations. Second, what can be done and what cannot be done have to be interpreted in terms of legal rules. This can be extended to general moral rules, but cannot be interpreted in terms of factual ability of agents. So, nobody has the RIGHT to rob you just because if you don’t give him the money, he “can” react by hurting you. Furthermore, also in case you actually have a RIGHT towards somebody, your reaction to non-compliance may not be “any” reaction, but it is further constrained. Third, in the common use of the terms that refer to the legal relations, they are not always ‘directed’ (Herrestad and Krogh 1995), i.e., they do not involve two agents, but just one. The correct interpretation of single-agent relations is obtained from DUTY and PRIVILEGE by quantifying over X: for DUTY, there exists an X that may react (or a set of X’s, having predefined features, e.g. policemen); for PRIVILEGE no such X does exist (notice that the negation has wider scope than the quantifier). What one should not say is “I have the right to do that”, since in this context the correct expression is “I have the privilege to do that”.

We can now return to our planning-intention perspective. From this perspective, Hohfeld’s approach tells us that, in case an action is a single actor one there are two possibilities: constraints and absence of constraints (DUTY and PRIVILEGE). But duties can also be ‘directed’ (and in such a case, they can be called ‘rights’ from the perspective of the other agent) and negated (no-rights): again, constraints and absence of constraints, although expressed in a slightly different way. It may be argued that privileges⁶ are useless (although it is useful to have a linguistic term for them, of course): if no constraint exists about an action, then the privilege is automatically obtained, without having to state it explicitly: privileges should not be part of legal (or moral) systems. However, they are useful for at least three reasons. First, they help in determining if an agent is liable for having done something: if the privilege is stated explicitly, then one does not need to search for all possible constraints related to the involved action. Second, rules are often organized in a taxonomy; so a rule could state that all agents of a given category (call it C) and being in a given situation type (S), have the duty to do A; but it may be the case that the rule does not apply to a specific subcategory of C (i.e., C’) or does not apply in a given subtype of S (i.e., S’) or both (*Lex Specialis*). So, one can say that, although, in general, the members of C have the duty to do A, the members of C’ (or the members of C, if they are in a situation of type S’) have the privilege not to do A. Note that we are not claiming this is the way an actual legal system works, but rather an idealization of it. In general there could not be an exhaustive coverage of all

situations, so that in real systems, privileges are even more important than in their idealized counterparts. The third, and perhaps the most important reason for the explicit statement of privileges depends on the hierarchical organization of the authorities able to enact laws: if a higher power wants to prevent that a lower power unduly constrains the behaviour of the entities subject to the law, it can state that these entities have some privileges: if the Constitution states that citizens have the privilege to express freely their political opinion, then no state or regional law can state the opposite (at least in principle).

With respect to the placement of the legal relations in the ontology, DOLCE states that they are subcategories of *Modal-Description* (which are *S-Descriptions*), and have as a component a *Functional-role* (e.g. an *Agent-role*), which is in a specific *Modal-relation* with some *Course* (i.e., a course of events). The modal relations are *has-right-on*, *has-not-right-on*, *has-duty-of*, *has-privilege-of*. All *Modal-Descriptions* are parts of an *S-Description* including two components: the first of them is a *Functional-role* and has as *Modality-Target* the second one (a *Course*). In other words, a *Duty*, for instance, is the description of a situation involving some action (a course of events is something that describes an action) which an agent (role) has the duty to perform.⁷ This approach seems to be logic-biased, where a *Modal-Description* recalls a formula in deontic logic. On the contrary, we are more agent-oriented, so that we need to associate with legal relations some more information, i.e., the author of the norm and the sanction. Also, the relation between *Norm* and *Plan* has to be made more explicit. Moreover, power-related relations (see Section 4) are also viewed as *Modal-Descriptions*, while it seems more reasonable, according to the usual analyses, to take them as relating an agent to another legal relation, which is the *p-effect* of the associated plan.

With respect to norms, in DOLCE, *Norm* is a subcategory of *Regulation*; *Norms* have the function of constraining *Perdurants*, and their relationship with legal relations is rather indirect. In our model, we assume that the legal relations are subcategories of *Norm* in that they specify (describe) particular constraints on behaviour (activities).

Let us now introduce an example. As we have seen, a *Plan* has *p-preconditions* and *p-effects*. Let's consider a plan schema. Its *p-preconditions* pose some constraints on the applicability of the plan. Therefore, the *p-preconditions* for cooking spaghetti with garlic, oil, and chili pepper (GO&CP), is to have spaghetti, oil, chili pepper, water (to boil the spaghetti), a pan, and some fire (a very easy recipe). As it is usual, *p-preconditions* are conditions that prevent the action from being executable (physical, or practical conditions). But it could happen that the King of That Kingdom has prohibited this recipe. Or, perhaps, this prohibition applies just to dinners with more than seven participants, or to dinners served after midnight, and

so on. Also, in some situations (here and now) garlic could not be a nice ingredient (because of its side-effects on your breath); or you know that one of your guests does not like chili pepper. All of these are not *p-preconditions*: if you have all the things you need, you can cook them in all situations described above (may be you are a spaghetti revolutionary, or you are not very polite). Anyway, these rules constrain your freedom to cook spaghetti GO&CP; you can pay for that: the King may imprison you, you may lose a kiss from your love (who does not appreciate, for this night, the smell of your mouth), or your guests may decide that they will never accept another invitation from you.

The “norms” we have exemplified above are not *p-preconditions*, but constraints on plans. We stress that in all of these examples, the utility of that recipe is reduced by the presence of the norm, so, it is possible that in absence of the norms, your intention for tonight dinner is to cook spaghetti GO & CP, but in presence of one or more of these norms, your intention is to prepare rice with mushrooms: you have chosen a different plan, whose utility is lower than the plan of spaghetti without norms, but is higher than the plan of spaghetti with norms (since the norm reduces the utility). The problem, now, is how to express the norms.

Let us consider the case of the King’s prohibition in its basic form (independent of the number of participants and of daytime). The rule involves three or four elements: the involved plan, the authority who posed the norm, and the agent who is taken as responsible for the execution of the action; the fourth element is not strictly required, but it is very common: it is the sanction⁸ the agent will undergo in case he does not respect the rule. So, a norm is not a plan, but ‘refers’ to a plan, in the sense that it (as a description) includes (is related to) a description of the action (i.e., the plan) which is forbidden (or permitted).

In Figure 5, we introduce the basic notions related to norms, and apply them to our spaghetti example. In this example, we introduce instances. They are represented as rectangles. However, see below for some comments on the distinction between classes and instances. Instances are linked to the category which they belong to via a thick dotted arrow. In the figure, it is stated that

1. *Norms* are a subcategory of *S-Description*. This should not be confused with the written law which is just a physical way to express the description: it is a medium for communicating the norm, not the norm itself. In fact, not all norms are written: they can be transmitted via verbal traditions and customs.
2. *Norms* have a *topic*,⁹ the *Plan* the norm is about.
3. *Norms* have a *polarity*. Its range is the abstract set composed of the two elements *pos* and *neg*. If the value of *polarity* is *pos*, then the norm




Figure 5. Norms and plans.

refers to something an agent must do, otherwise (as in our example), it refers to something he must not do. In other words, if the polarity is *pos*, an agent *X* has the *DUTY* to do the action; if it is *neg*, it has the *DUTY* not to perform the action.¹⁰

4. Every *Norm* has a *bearer*. More precisely, it has a set of bearers: usually, norms apply to categories of agents, and not to single individuals. As usual, a *relation* arc enters another concept: that concept is associated to the set of potential bearers of the given norm.
5. Every *Plan* has an *actor*, who is the ‘main character’ of the plan. Although the *actor* of the plan is conceptually distinct from the planning agent, we assume, in the current model, that the two roles are played by the same individual (see Figure 4, and the comments in the text).¹¹ In the present example, we assume that the host of the dinner is the *actor* of the plan of serving GO&CP; so, he also is the agent that must decide about the recipe for the dinner (the planning agent).
6. The *actor* of a plan, if that plan happens to be the *topic* of the *Norm*, and if that *actor* is of the right category, becomes an (individual) *bearer* of the rule. Of course, this applies also during the planning phase, so that the planning agent, which is a potential *bearer*, may discard the plan not to become an actual *bearer*. The identification of the actor with the bearer may appear as an oversimplification, but see (Kralingen et al. 1999, p. 1136): “The fourth slot is reserved for the agent who performs the act (the agent and the norm subject must be the same person)”.
7. Every *Norm* is associated with a *sanction*,¹² which is a *PlanWith Patient*, i.e., belongs to that subclass of *Plan* which is in some *has-patient* relation with an individual playing a *Patient-Role*.

8. The *bearer* of the *Norm* is the *has-patient* of its *sanction*. Of course, this does not take into account the notion of *responsibility*, but we assume that being an agent who freely chose a given line of behaviour is the basis for further elaborations of this concept.
9. The *author* of a *norm* is an *Authority* (see Section 6).

With respect to the distinction between categories and instances, a few words are in order. The intuition is rather simple, and well known from decades of research on AI and KR: in standard model theory there are individuals and relations; an instance is an individual of the model, while a category corresponds to a one-place relation (expressed in logics as a one-place predicate). So, we have to define the individuals in our model. Clearly enough, the King is an individual (in particular, an *Endurant*), as well as the abstract entity *neg*. What about norms and plans? Is the single norm *NoSpaghettiGO&CP* an individual? It depends on the way we use it in a reasoning framework. If we have to discuss a case, then we have to apply the norm to the description of the case; in doing that, we must bind the variables (roles) associated with the norm; in particular, we must bind the bearer. In doing that we obtain an entity which is more specific than *NoSpaghettiGO&CP*, but seems to be of the same ‘type’. So, it seems more appropriate that *NoSpaghettiGO&CP* be a category, whose instances are the various applications of the norm to real cases. However, if we refer to the norm as a description, then *NoSpaghettiGO&CP* is an individual, i.e., a single individual description, which may be in a *refines* relation to another individual norm description. The case of plans is even more striking: is the plan *GO&CP* an individual or a set, i.e., the set of all executions of the plan? The solution we propose is the same as above, but here we have three levels, instead of two: there are various descriptions of the *GO&CP* plan (instances of *Plan*); there is the execution of the plan (an instance of a *Perdurant*); but we also have the particular plan the planning agent has chosen (its *Intention*, which is a *Complex-Task*; see Figure 4).

4. Authorities: the legal relations concerning *Power*

Traditionally, law scholars like (Hart 1961) distinguish between primary laws, whose purpose is to direct the behaviour of citizens, and secondary laws, which, among other functions, serve to the maintenance and dynamic management of the normative system. These rules form a “subsystem of rules for change”: rules which have legal effects and which are instrumental to the primary system, in that they regulate the regulation [e.g., art. 2 of Italian Civil Code: “the creation of laws [...] is regulated by constitutional laws (Cost. 70)”]. This subsystem, according to Hart, does not include only

the rules of change which specify how new laws are introduced or old ones removed, but it also includes rules about “powers for the private citizen”. These rules are at the basis of civil code and regulate testaments and contracts; for Hart they allow the exercise of limited legislative powers by the citizens. These rules do not create or remove general laws but they introduce and remove individual obligations and permissions: e.g., in the Italian Civil Code, art. 1173 (sources of obligations) specifies that obligations are created by contracts (a contract being an agreement among two or more parties to regulate a legal relationship about valuables by art. 1321).

In Hohfeld’s framework, this distinction is reflected by the introduction, beyond the four Legal Relations we discussed in Section 3 (RIGHT, DUTY, NO-RIGHT, PRIVILEGE), of the relations of POWER, INABILITY, LIABILITY, IMMUNITY. They concern actions affecting the existing Legal Relations. In our terms, the topic of this particular type of rules is that subcategory of *Plan* whose instances have, among their effects, a modification of a Legal Relation. As we have seen in Figure 5, *Duty* is viewed as a special subcategory of *Norm*. Accordingly, it seems reasonable to assume that also *Powers* are particular subcategories of *Norm*. So, we get:

1. *Duty, Right, No-Right, Privilege, Power, Inability, Liability, Immunity* are specific subcategories of *Norm* (just *Duty, Privilege, and Power* are shown in the figure).
2. Some *Plans* have as their effect *ConstraintsOnBehavior*, which is a specific subtype of *S-Description*. They are *JurPlans* and *LawPlans*.¹³ The difference is that *JurPlans* establish new *IndividualConstraintsOnBehavior*, while *LawPlans* establish new *Norms*. An example of *JurPlan* is ‘renting a car’, where, if the plan is executed correctly and by a person who can do it (according to an existing *BasicPower* norm), new *IndividualConstraintsOnBehavior* are set up, regulating the behaviour of the two parties (duties and privileges of the renter and of the client). On the contrary, the introduction of a new norm stating that car rental may be made just to people older than 20 is the execution of a *Law-Plan*.
3. Not every agent may execute *JurPlans* or *LawPlans*. An agent who can do that is said to have *Power*. For example, the owner of a car rental agency has the *Power* to rent a car, and, according to the hypothetical rule mentioned above, anybody older than 20 (possibly, with a driving license) has the *Power* to hire a car from an agency. Both of these *Powers* are *BasicPowers*, since they enable the execution of plans that establish individual duties. On the contrary, only people (or institutions) having some *LawPower* can modify the minimal age for renting a car, since this implies a modification of the legal system.




Figure 6. Norms concerning Power.

The result of this conceptual analysis is depicted in Figure 6. The ontological modelling of point 3 above does not appear in the figure, but it is (partially) reported in Figure 7.

We can now turn to the four Legal Relations concerning *Power*. Their interconnections seem less clear than the ones among non-powers.

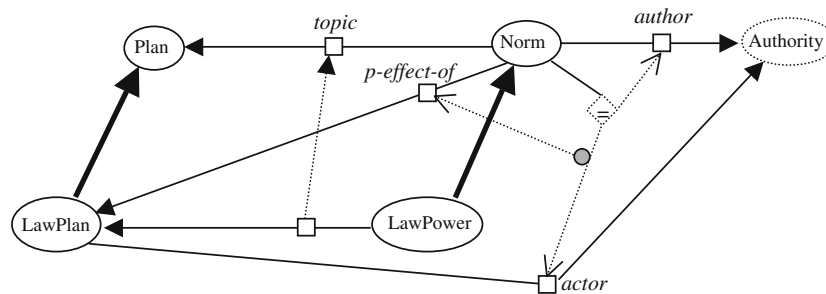


Figure 7. Authorities: what they do: the *author* of the *Norm* is the same as the *actor* of the *LawPlan* that brought about (*effect-of*) the *Norm*. This is obtained by exercising some *LawPower*, which has as *topic* the executed *LawPlan*. Note that there are two *Norms* involved here: the first one is the *LawPower* (a specific type of *Norm*) which has been executed; the second one is the ‘new’ *Norm* that has been created.

1. X has the **POWER** to bring about that P (where P is some Legal Relation), iff there exists an action A such that in case X does perform A, P is established. In other words A has among its effects P.
2. X has the **INABILITY** to bring about that P (where P is some Legal Relation), iff there exists no action A such that in case X does perform A, P is established.
3. P has the **LIABILITY** of being brought about by X iff there exists an action A such that in case X does perform A, P is established.
4. P has the **IMMUNITY** of being brought about by X iff there exists no action A such that in case X does perform A, P is established.

Although the picture is simpler than in case of non-power relations (2 is the ‘inverse’ of 1, 3 is the ‘opposite’ of 1, and 4 is the ‘opposite-inverse’ of 1), there is just one agent involved, and the terms **LIABILITY** and **IMMUNITY** apply to Legal Relations, while in ordinary language they are more commonly used to refer to agents. When one says that agent Y is ‘liable’, it must be assumed that the speaker refers to a legal relation P involving Y, so that ‘Y is liable to P’ is interpreted as ‘there is some agent X and some action Act, such that if X performs Act, then a Legal Relation (P) involving Y is established’ (or, alternatively, Y is liable to P iff there exist an agent X such that P has the **LIABILITY** of being brought about by X; remember that P ‘involves’ Y).

Consequently, since all other relations are easily defined in terms of **POWER**, we can now focus on it. In principle, *Power* does not tell us very much: there is some action that an agent can perform that has some effect concerning Legal Relations. But, apart from the final qualification about the type of effect, this is true for every action: this is the very definition of an action (or a plan, if the action is complex): a plan is something an agent may execute to achieve some effect.

So, it is sufficient that such an action be defined, in order to obtain *Power*: if the action does not exist, then no *Power* exists; if the action does exist, then there is the *Power*; but provided that the preconditions are satisfied! And the origin of *Power* seems to be in the preconditions. As we have already noticed, usually, preconditions express factual constraints for the execution of an action: if there is no lamp, you cannot turn on the light. But in this case, something different is involved: there are some preconditions that are ‘conventional’, in the sense that even if they do not hold, the action appears to be executable. So, if a judge looks at a person while walking in the street, and tells him: “I condemn you to two years”, the ‘condemn’ action seems to have been executed, but, of course, this is not so. But where do these further constraints come from? The answer is that there needs to be some *Norm* (or a set of norms) that defines the action. So, a *Power* depends on the existence of norms defining an action modifying a Legal Relation. Under this reading, the

Hohfeldian definition applies without changes: what is important is just to note that the action A, whose execution can bring about P, is not a ‘natural’ action, but it is a ‘conventional’ action suitably defined.

We must remark that what is in order here is not the introduction of a *Norm* (an action which always require an *Authority*), but the introduction of a *Plan*, as, for instance, the ‘selling and buying’ plan. The fact is that, according to the previous discussion, there is no way to define this plan, other than having a (set of) norm introducing it, i.e., no *JurPlan* or *LawPlan* can exist, unless it is the topic of some *Norm*. In other words, an agent cannot ‘invent’ the plan ‘selling and buying’, as he could ‘invent’ a new plan for organizing a party: the latter can be invented, because its effects are factual effects (if the party comes out to be a nice party, the plan is a good plan, since it has achieved the effects it was invented for); but the former cannot be invented, since it will never have any effect at all, unless some authority establishes which are these effects, and under which conditions they are achieved (see Figure 6).

For instance (see Figure 6), a *LawPlan* for the definition of new types of contracts can be executed for defining *UnsafeContract*. This execution has as its effect the creation of a *Norm* (the one ruling *UnsafeContract*). This is a *Power* (the one that is possessed by those who can use *UnsafeContract*) i.e., a type of *Norm*, but it does not seem that there is a newly defined *Plan* involved. However, each *Norm* refers (*topic*) to a *Plan*; so, the same holds for *Power*; in our example, it is the *Plan* that must be executed for entering a *UnsafeContract*. In general, it is not possible to introduce a *Norm* without the associated *topic*, and the *topic* of a *Power* is the *Plan* that must be executed to exercise the *Power*. Note also that we are referring to the *creation* of norms, but of course nothing changes substantially in case of *modification* or *cancellation* of norms.

The definition of *UnsafeContract* was an example of application of *Law-Power*. On the contrary, the execution of the *Plan* associated with *UnsafeContract* is just an application of a *JurPower*. In this case, what come to existence are just one or more *instances* of Legal Relations (e.g. new *Duties* and *Privileges*). It is clear that the first type of *Power* has a much greater impact on the society, so that it must be much more constrained. So, the *Authority* which can do that is restricted to be a carefully defined group of people (or the King, in our spaghetti example). Of course, the degree of impact is different for different laws, so that different levels of authority are usually defined.

The task of determining “who has power” is much more complex than it might appear. Each and every institutional organization (ranging from the Constitutional Court to the lowest levels) has received its POWER on the basis of some set of norms, stating which are the attributions of these organizations and how they are composed. So, they find their proper place

within an established social environment. But what about illegal organizations, for instance Mafia? How can we model the fact that a local Boss has established that all food shops of a given area must pay 1000 Euro per month for the protection? It seems that this is not so different, from an ontological point of view: even in this case, a rule has been introduced, together with its restrictions (“food” shops, “a given area”), and plan (pay 1000 Euro), in this example with a POS polarity. So, there should have been an exercise of POWER; but which is the source of this POWER? It may be argued that it is force and violence, but perhaps it is more precise to say that it comes from a parallel social organization whose grounding is on force and violence. Of course, the fact that, hopefully, there will be greater resistance to this rule than to legal rules, depends on an evaluation of the source of the POWER (commonly agreed Constitutional origin rather than violence), but this does not affect the ontological status of the rule itself. So, we have that *Authority* (i.e., the owner of POWER) gets its status from a social body, a source which is something different from the *Authority* itself.

With respect to *BasicPower*, it must be added that all agents are full of *BasicPower*. As legal persons, we are owners of some properties, so that we can do things with them, as selling them, whose effects are guaranteed by the legal system, or more precisely, by the *BasicPower* it assigns to us (unfortunately just in most cases, not always, see Sen (1999) and the concept of *capabilities* he introduces). Now, we can decide that *Authorities* are agents having a *Power*, in which case all of us are *Authorities* with respect to our home or our credit card, or that the term *Authority* applies just to agents having some *LawPower*. This latter usage seems more akin to the standard one, so we adopt it.

Before turning back to the technical issues, we must observe that some rules exist independently on any POWER. They are the rules of tradition (not to be confused with the rules coming from written religious texts). They are rules the origin of which is unclear: rules to which everybody in a society usually conforms, just because they exist, and because who does not respect them is not respected by other people. This situation does not seem to pose any special problems for the ontological framework we are proposing: the basic structure of the norm is exactly the same as before. Simply, the author of the rule is undefined. But a sanction usually exists, and it is the probable reaction of the other members of the society (in extreme cases, nobody wants to have any contact with you any more).

Now, we show the placement of *Authority* in the ontology. Note that *Authority* is a material role, since individuals or groups can become or cease to be authorities, according to the application of some POWER. The main ideas are depicted in Figures 7 and 8.

In Figure 7, we see that the *Author* of a *Norm* is the *Actor* of the *LawPlan* whose effect is that *Norm*. In Figure 8, it is shown how one becomes an *Authority*, i.e., as an effect of a special type of *LawPlan* named

attributions (*Privileges, Duties*, and so on). This is a true application of *LawPower*, a task of some pre-defined *Authority*.¹⁴

5. A case study: *fruits*

In this section, we report the result of a preliminary study on a legal concept, i.e., fruits (see Italian Civil Code, art. 820). A *Fruit* is something that is obtained by somebody as the result of his ownership of something else. For instance, you obtain the apples if you are the owner of an apple tree, or you obtain some money if you own a house and you rent it out. Notice that the presence of fruits does not affect the ownership of the thing. This section aims at presenting knowledge about fruits appearing in the Italian Civil Code. In the next section, we will report the representation of a norm concerning fruits, in order to get a first rough idea about the impact of norms on the ontology.

5.1. GOODS

Goods are defined as a subcategory of *Endurants*¹⁵ (see Figure 9). We partition the goods in two subcategories: the goods that may undergo a property relation (*SubjectToProperty*), and the ones which cannot (*NotSubjectToProperty*). The individuals of the first category can play two roles: goods that, at some time, are a property of somebody (*Asset*), and goods which are temporarily not owned by anybody (*ResNullius*). The *PropertyRel* relation involves a *Legally-Constructed-Person* (the *owner*) and a *SubjectToProperty* good (the owned thing: *ownee*). An *Asset* is the *object* of a *SubjectToProperty* good, for which there is a *PropertyRel* relation (i.e., the *prop-of* role is restricted to have at least one filler: 1,n). On the contrary, a *ResNullius* is not involved in any *PropertyRel*: *prop-of* is restricted to have no filler (0,0) (Figure 9).

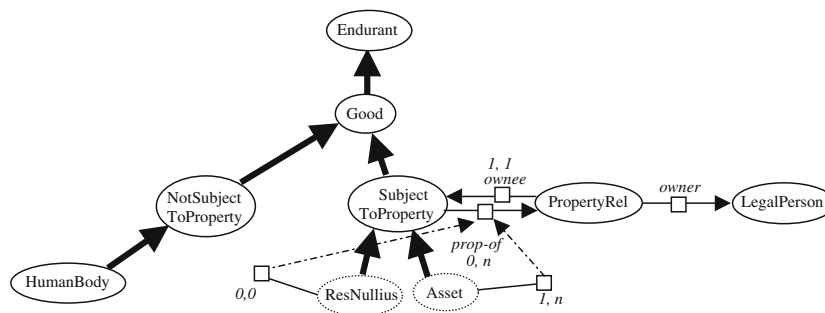


Figure 9. Goods.

5.2. FRUITS

Fruits are *Goods*. There are two types of *Fruits*: *NaturalFruits* and *CivilFruits*. *NaturalFruits* are goods that are “detachable” parts of some “original” good, and which originated out of the original good via some natural process (e.g. the growth of real fruits – apples, peaches, etc. – the birth of a lamb out of its mother, etc.). In this paper, we will not discuss natural fruits any more.

On the contrary, *CivilFruits* are *Assets*, which have changed their *owner*, after some more or less complex procedure. The general idea is that some agent *A1* owns (and possesses) some *Good G1*; then, he can decide to exchange the possession (maintaining the ownership) against some *Good G2*, usually for some limited period of time. A second agent *A2* gets the possession of *G1* and gives *G2* to the first agent. In this situation, *G2* is a *CivilFruit* of *G1*. The simplest example is the one of (house) rental, where *G1* is the rented house, *A1* is the owner of the house, *A2* is the tenant and *G2* (the *CivilFruit*) is the amount paid for the rent. All of this is ruled by some *JurPlan* (as a *Contract*), which must be executed by the involved parties (*A1* and *A2*). As seen before, in order to be authorized to execute a *JurPlan*, an agent must have some *JurPower*. In the example of house rental, this power has been given to the owner *A1* by the original act of buying the house. On the contrary, the power exercised by *A2* is due just to his being a *Legally-Constructed-Person*, and to his power on the money he pays for the rental. Of course, all of this produces new legal relations, in particular, new *Privileges* for *A2*; these are the *Privileges* associated with the *Possession* relation, which has been created as an effect of the execution of the *Rent* contract.

In Figure 10, we have depicted some of the concepts discussed above. In particular, it has been made explicit that:

1. *CivilFruit*, since it is an *Asset*, has a *owner* (linked to it via the *PropertyRel* concept).
2. *CivilFruits* come from some other *Good*, which we have called *ExploitedGood* (which is a role taken by the *Good* when it is used in a *FruitProducingContract*; see point 10 below).
3. The *owner* of the *ExploitedGood* is the same as the *owner* of the *Civil-Fruit*.
4. There exists a special type of *Contract* involving exactly two contractors (*Contract#2*).
5. A subtype of type of *Contract#2* is *FruitProducingContract*. The two contractors are called the *giver* and the *receiver*.
6. *FruitProducingContract* involves two *Goods*, called the *primary* and the *secondary* (in this case, the names are not so clear; we had to keep apart the good which change its possession, as the rented house – the *primary* – from the good that changes ownership – the *secondary*).

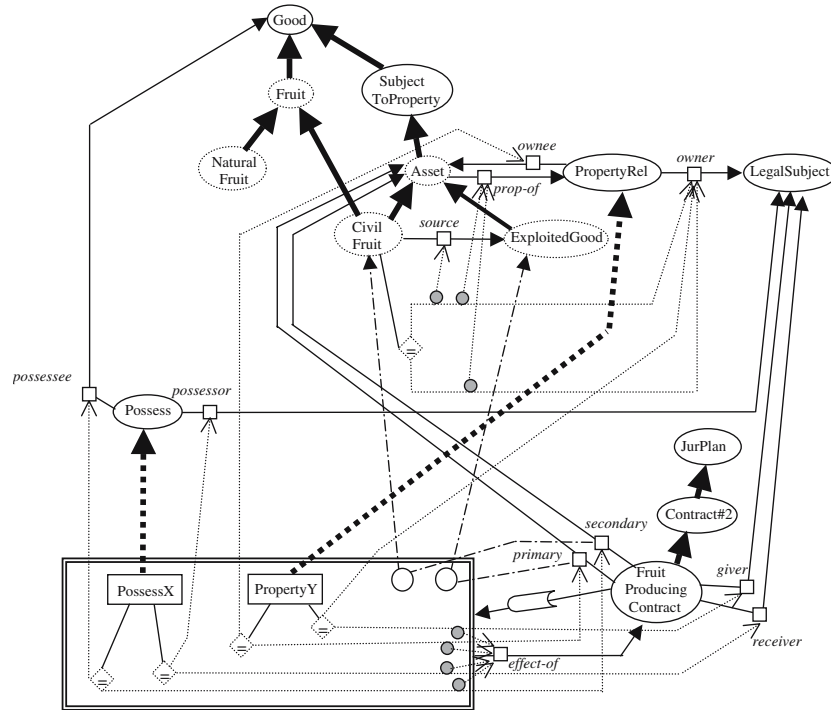


Figure 10. Fruits.

7. A first *effect* of a *FruitProducingContract* is that a new instance of the *Possess* relation is created (*PossessX* in the figure), whose *possessor* is the *receiver* of the *FruitProducingContract*, and whose *possessee* is its *primary*.
8. A second *effect* of a *FruitProducingContract* is that a new instance of *PropertyRel* (*PropertyY*) is established.
9. The *owner* of *PropertyY* is the same as the *giver* of the *FruitProducingContract* which produced the creation of *PropertyY*, while the *ownee* of *PropertyY* is the *secondary* of *FruitProducingContract*.
10. The *primary* of *FruitProducingContract*, which is an *Asset*, as an effect of an execution of *FruitProducingContract* takes the material role of *ExploitedGood*, as well as the *secondary* takes the role of *CivilFruit*.

It is clear that this is just a simple example. However, it makes explicit some interesting features; in particular, items 2 and 10 above. Item 2 implies that there is no *CivilFruit* without an *ExploitedGood*; item 10 specifies how *CivilFruits* come into being; note that this is not a definition, unless one states that the *only* way for an *Asset* to become a *CivilFruit* is as an effect of *FruitProducingContract* plan, which is not stated here. But apart from

contracts other than *FruitProducingContract*, what can easily be made is to introduce *specializations* of *FruitProducingContract*. So, we can have *ApartmentRental*, *CarRental*, *LandRental*, *BankDeposit*, etc. all of which share (inherit) the general features of *FruitProducingContract*. In particular, all of them have in common the effect that some *Asset* becomes a *CivilFruit*.

5.3. A NORM ON FRUITS

As an example of norms concerning fruits, we consider part of the article 1477 of the Italian Civil Code: the seller must give the buyer the possible fruits produced by the sold asset (the *ExploitedGood*). More precisely, if the *owner* of a *Good*, which is currently playing the role of *ExploitedGood* sells that *Good*, then also the associated *Fruits* must be given to the buyer.

In Figure 11 we show the relevant part of the representation. The focus of the example is the *Norm* instance *Give-FruitW*, which is a *Duty* that comes into being as an effect of the execution of *SellExploitedGood*, which is a *JurPlan*. This instance is analogous to the instance of *Property* that refers to the fact that the buyer is now the owner of the sold item, in the sense that both of them are effects of a selling action. Note also that the *Norm* mentioned in the previous paragraph is not represented explicitly, since it is actually part of the definition of the *Power* associated with the definition of *SellExploitedGood*. In fact, as stated in Section 4, the issuer of the Italian c.c. article 1477 had the power to (implicitly) define the *SellExploitedGood* action (although no name was given to it). Within this definition, the effect of




Figure 11. A norm about selling *ExploitedGoods*.

creating the *Duty* of transferring the *CivilFruits* was included, and this is what is depicted in the figure.

1. *Selling* is a specific type of *Contract #2* (and hence a *JurPlan*).
2. Among the effects produced by *Selling* there is a change of *Property*: a new instance of *Property* is established (shown in the figure) as well a previous property is cancelled (not shown in the figure).
3. In case the sold *item* is an *ExploitedGood*, a specific subcategory of *Sell* is defined: *SellExploitedGood*.
4. Beyond the (inherited) effects produced by *Sell*, *SellExploitedGood* produces a further *effect*: a new *Duty* (*Give-fruitW*). The *bearer* is the same legal subject who plays the role of *seller* of the *Selling* action, but this need not be specified explicitly, since it is inferred from the general rule that the *bearer* is the *has-agent* of the *topic* of the *Norm* (i.e., the *giver*).
5. On the contrary, it is explicitly stated that the *giver* of the *Civil-Fruit* is the same as the *seller* of the *ExploitedGood*.
6. The *Assets* which must be given by the *seller* to the *buyer* are the *Civil-Fruits* that are produced by the asset which has been sold.
7. According to the Italian civil code, when a duty created by an obligation is violated by an agent, he must repay for the damage created to the beneficiary: the sanction associated to the *Give-fruit* duty is thus to repay the buyer for the damage due to not giving the fruits of the sold assets.

6. Ontological models of the legal system

Classical formalizations of legal reasoning are based on deontic logic, as von Wright (1951), Alchourron (1993), van der Torre and Tan (1995), Prakken and Sergot (1997), to quote just a few examples. More recently, ontologies have been adopted as a means to give structure to the legal domain, beyond the basic deontic primitives. Ontologies have been used in a wide range of applications, from legal advice [see the application of the CLIME ontology to Maritime Law (Winkels et al. 1999; Boer et al. 2001), the representation of cadastral data and norms ruling the Real Property Transactions (Stuckenschmidt et al. 2001) and the reasoning about the eligibility to legal aid for the indigent (Zeleznikow and Stranieri 2001)], to the access to legal information (Winkels et al. 2000; Tiscornia 2001). Various studies aimed at proposing ontological primitives (Mommers, 2001) or at comparing and evaluating different approaches to legal ontologies (Visser and Bench-Capon 1998; Bench-Capon 2001). In this section, we overview three models, which seem more relevant for setting the context of the present work.

In various papers, Breuker, Valente and Winkels (BVW) have described a 'functional model' of legal systems. The model, which is very articulated, and covers many aspects of legal knowledge and reasoning, is based on the idea that

“the legal system as a whole (and therefore each of its components) exists to accomplish a certain function, in order to obtain certain social goals. The legal system is thus viewed as an entity with a certain internal structure, behaving in an environment, and that was designed to work in a certain way in order to be able to accomplish specified social goals. The main function of the legal system is to change or influence society in specific directions, determined by certain social goals” (Breuker et al. 1997).

The goal of BVW's research activity is to provide a general framework for law, including the legal reasoning necessary to assess the cases. The knowledge required has been classified according to its function: *Normative knowledge* (it prescribes behaviour and defines a standard of comparison for the social reality), *Meta-legal knowledge* (it organizes the relative position of norms, coping with conflicts, validity, etc.), *World knowledge* (it provides the background on which norms are built and deals with both entities and behaviours), *Responsibility knowledge* (it refers to the links between causal relations and legal responsibility of agents), *Reactive knowledge* (how the legal system reacts to the behaviour of agents; i.e., with sanctions or rewards), and *Creative knowledge* (enabling the system to define new entities and new types of entities). A general schema of legal reasoning involves taking in input a case, 'translating' its description in legal terms (i.e., according to legal categories), then 'matching' it against the existing norms in order to verify if there has been a violation, then assessing the responsibilities of agents and possibly choosing the reactions of the legal system.

It is clear that BVW's perspective is much wider than ours. The present article, in fact, contributes mainly to the definition of the links between *Normative knowledge* and *World knowledge* (with a possible future impact on *Responsibility knowledge*). What is especially relevant in the present context is BVW's view of model-based reasoning; their approach moves away from pure heuristic reasoning towards a full representation of the principles and contents of legal knowledge. But they claim that full model-based reasoning is not adequate for the task at hand, so that "weaker versions of model-based reasoning that do not require full envisioning" (Valente et al. 1999, p. 1082) are preferable.¹⁶ It cannot be denied that full-blown model based reasoning is hardly feasible in a domain as complex as legal reasoning. However, it seems that some further steps towards this end can be useful to approach the goal of linking legal concepts to world knowledge.

In particular, we follow the suggestion of matching real-world situations against legal knowledge, but we also claim that this match can produce deeper insights on what happened and why, if the mental processes that lead people to behave in a given way are properly taken into account. Also in this respect we follow BVW, who state that:

“Primary norms address individual behavior (instances of behaviour) in the world by means of reference to patterns of behavior” (Valente et al. 1999, p. 1093).

and

“... in law the causal reasoning is contingent upon notions of responsibility, i.e., one has to take into account also the concepts ‘intention’ and ‘knowledge/belief’” (Valente et al. 1999, p. 1102).

Our view is that many of the concepts strictly linked to the notion of agency (i.e., behaviour, intention, knowledge/belief; see the research on BDI agents) are relevant in legal modelling, but they still lack an adequate formalization in terms of ontological primitives. We claim that BDI models, together with their involvement of desires and intentions, are one of the essential bricks required to expand and extend BVW’s analysis. In fact, in assessing a case, it is not only necessary to decide if one or more norm have been broken, but it is also necessary to understand the reasons why an agent did break them. And this cannot be done, unless one has a view of the reasoning mechanism that led the agent to choose that specific line of behaviour. This is also made clear in the description of *Agent causation* in (Lehmann et al. 2003):

“The agreement around cases of agent causation is not reached as easily as in case of physical causation. This is due to the problem of detecting the beliefs, desires and intentions of the agent that starts the physical process. Usually a non intentional action does not fully qualify the agent that performs it as an agential cause, despite him being the physical cause” (Lehmann et al. 2003, pp. 21–22).

The present article tries to specify which is the link between an agent’s intentions, her/his behaviour and the norms that affect them. In this sense, it can be seen as an extension of BVW’s model, although our focus is more on the reasoning of the agent than on the reasoning of the legal system that has the task of inferring (abducing) it from the agent’s behaviour.

Similar comments apply to the system described in (Kralingen et al. 1999), where the “conceptual frame-based ontology” includes three types of frames:

norm frames, act frames, and concept-description frames. In the description of the Norm Frames, the authors state:

“A legal subject attempts to stay within the law and attempts to avoid the breach of norms. The reason for this behaviour, whether it is social pressure or threat of sanctions, is not important to our discussion of norms” (Kralingen et al. 1999, p. 1133).

And, in fact, the norm frames include five slots, which describe the subject of the norm, the legal modality of the norm (duty-imposing, permission or power-conferring), the action to which the norm refers to (What must be done or forborne?), and some extra conditions of application (place and time). As we have seen, we assume that a sixth slot is required, i.e., the sanction. In our view, this is essential in the task of explaining how the existence of a norm affects the behaviour of agents, who have to take into account the various aspect of the situation (including the existence of norms and possible sanctions) in order to choose a line of behaviour. Actually, the act frame includes three slots that are related to agent behaviour: *cause*, *aim*, and *intentionality*, but it is not clear if they have any impact on the actual model. In any case, the glosses reveal that these terms are used in a rather different way than we do (see Section 4).

In order to give an example of the importance of a deep model of actions, we may refer to the domain used in (Kralingen et al. 1999), i.e., the Imperial College Library Regulation (ICLR). One of the norms refers to the maximum number of books that can be borrowed. In particular, it states that that number cannot be exceeded; it is split into two norms, the first ones specifies that a borrower cannot take a book, if s/he has already reached the limit; the second one says that the librarian should not lend a book to a person who has reached the limit. Since the system has the goal of assessing if some norm has been broken, this means that, in fact, norms can be broken. But why? If this can happen, then the librarian must have had some reasons for doing that (an error or a principled choice). In our view, the librarian must also have some reason for not always breaking the rule (sanctions); but if he did that, it could depend on the fact that the borrower was told by her thesis advisor that she had to read a book before the day after, that she has reached the limit, that she lives very far from the library, and that she promised to bring back one or more book the day after. In such a case, we assume that the librarian made the right choice, even if it goes against the norm. And this must be taken into account when the case is assessed, if any action has to be taken against him.

A third model which is relevant here is the one presented in Hage and Verheij (1999). It is based on three conceptual primitives, i.e., *states of affairs*, *events* and *rules*. States of affairs concern any type of piece of reality (as ‘John is a thief’, or ‘A minor cannot make a valid will’), and are changed by events

(as ‘John taking away the car of Gerald’ or ‘An international treaty being ratified’). There are subcategories of events:

“A special kind of events are acts: events that consist of the intentional behaviour of an individual. A special category of acts are the so-called juristic acts” (Hage and Verheij, 1999, p. 1049).

In our view, one of the most important features of Hage and Verheij proposal is the parallel existing between the pair <event,causality> and the pair <rule,constitution>. The events change the state of affair ‘obtaining’ (i.e holding) before the event into a new state of affair obtaining after the event. And the new state of affairs is causally related (via its ‘effects’) to the event. On the contrary, two state of affairs (that occur simultaneously) are related via a constitution relation (called supervenience) if there is a rule establishing the existence of this relation. For instance, the event ‘signing of a contract’ is causally related to the state of affair ‘the contract is signed’, and ‘the two contractors are under a contractual bond’ supervenes to the latter. It seems that the concept of supervenience is strictly related to what Pollack (1990) called ‘generation’: in her case, the difference is mainly a matter of description, one of the example being ‘press the switch’ and ‘turn the light on’: they are not different actions (the agent has not done two different things), but the first only produces the second if some conditions apply (there is electric power, the wires are not cut, etc.). In the case of supervenience, we have that the conditions for the existence of the relation are not physical, but conventional: it is not different to have the contract signed and to have two individuals under a contractual bond, provided that there is a rule stating the correspondence; so, it is mainly a problem of description of a state of affairs which actually is a single one. The ‘conventionality’ of ‘supervenience’ is of paramount importance, and has been discussed in Section 4. What is also interesting is that the chain of (inverse) supervenience starting from the signing of a contract passes through ‘one of the two contractors (A) is under an obligation towards the other (B)’, and ends in ‘A ought to perform some action’: again an explicit reference to agent behaviour.

According to the previous discussion, we claim that actions, the way people use their knowledge about actions, and the way people decide which goal(s) to pursue via their actions, are all elements that have a basic role in defining a legal ontology. The next section is devoted to the introduction of a few proposals about the representation of plans in ontologies.

7. Ontological models of plans

Although planning is one of the main topics of AI research, not many attempts have been made to provide a principled formalization of planning

concepts. Presumably, one of the main difficulties stands on the fact that knowledge about planning is mostly embodied in procedures, which are notoriously hard to represent declaratively (as an ontology demands¹⁷).

The input to a planning procedure (planner) includes at least the following elements: a set of goals of the planning agent, an expressions of the preferences of the agent, a description of the current situation, and background knowledge. The latter, in turns, includes both knowledge about entities and relations existing in the world, and knowledge about the actions that can be performed in the world (with conditions for their execution and, possibly non-deterministic, effects). The output of a planning procedure is an intention, i.e., a sequence (more or less detailed) of actions that the agent has decided to execute; so, an intention is a specific plan of actions which the agent has committed to. In principle, the planner has to produce all possible plans that enable the agent to achieve one or more of his goals (using the goals, the current situation and the background knowledge), and then must choose one of them (using the preferences), which becomes the agent intention. Of course, the process cannot be carried out in this way, because of the combinatorial explosion in the number of candidate plans; so actual planners become rather complex, because of the need to reduce the search space via heuristics and interleaving of various steps.

The most extensive attempt to formalize the notion of plan has been made within the Planning Initiative of DARPA, resulting in Shared Planning and Activity Representation (SPAR) (Polyak and Tate 1998, 1999). The results of this study are impressive, and their main features are the coverage of the formalization (which includes also probabilities and uncertainty) and range of sources which have contributed to the final model, among which representation languages [as PIF: Lee et al. (1998)], tools [as Prodigy: Carbonell et al. (1991)] and true ontologies [as Ozone: Smith and Becker (1997)]; see Polyak and Tate (1998, pp. 3–4) for a complete listing. However, the strength of SPAR seems to bring with itself some weaknesses: the fact that various requirements come mainly from industrial environments (e.g. the need for standards in process specifications), and the necessity to quickly approach the core of planning has made the top-level ontology somewhat confused, including as sisters (under the topmost concept ‘entity’) concepts rather dishomogenous, as *Environment*, *Activity*, *TimePoint*, *Issue*, *EvaluationCriterion*, *WorldModel*, *WorldModelSpecification*, to quote just some of them. It seems clear that many of the concept and features appearing in the SPAR proposal are of paramount importance; but it seems equally clear that the connections to a foundational ontology need be refined.

A model which suffers from some of the drawbacks of SPAR (i.e., unclear connections to a top-level ontology) is PLANET (Gil and Blythe 2000). However, it includes a very nice feature, i.e., an explicit representation of the partial structures which must be built and managed during the planning

process. PLANET includes four main structures: *planning-problem-contexts*, which describe the background knowledge, the constraints that the resulting plans must satisfy, and the initial situation, *plan-task-descriptions*, which refers to plan templates (i.e., knowledge about actions) and plan tasks (i.e., instantiated portions of plans), *goal-specifications*, which describes the goals the planner must try to achieve, and *plans*. We want to stress here just that there is a relation between plans and *planning-problem-context* called *candidate-plans*, which carries information about the suitability of the plan for that problem; the possible values of this feature are *unexplored*, *rejected*, *feasible*, *selected*. So, it is modelled the fact that plans are built as solution hypotheses for a problem, and are then evaluated, on the basis of some criteria, in order to make a final choice. This is especially important in our model too, where norms have exactly the role of further constraints on plans, affecting the final choice of the agent.

8. Conclusions

In this article, we have presented an ontological model of norms, based on the behaviour of agents. We have shown that norms can be characterized as *constraints on behaviour*, i.e., as statements specifying what an agent can and cannot do, and what happens when a norm is broken. The basic idea is that agents choose a line of behaviour according to the utility they may gain; and usually the breach of a norm produces a decrease in the estimated utility, because of the risk of being sanctioned.

In order to include in the ontology knowledge about behaviour, we have provided a sketch of an ontological representation of plans and intentions, showing how a *Norm* can affect the planning process: the presence of a *Norm* can urge an agent to adopt a given line of behaviour (complying with the norm) instead of another (breaching the norm). This is obtained by introducing the notion of utility and by letting the agent determine his intention according to the expected maximal utility.

Although the article has mainly addressed primary norms, by also showing how they are related to the basic Hohfeldian Legal Relations, it has also shown that the proposed model covers *authorities*, i.e., it enables one to express explicitly the role of authorities in a legal system, and the way they are created. Finally, we have shown how the model applies to a specific case study (*Goods and Fruits*).

Two final comments concern the degree of coverage of the article with respect to the various components of a legal system and the complexity of the representation. As we have stated in the introduction, the goal of this paper is not to develop a general model of a legal system (see the models described in Section 6), but only to give an ontological characterization of a specific

aspect, i.e., the role of primary norms. In principle, norms are just a small portion of such a model, although they are a very important component, as we have argued in Section 2. Representing norms in an ontologically plausible way is the first step towards the implementation of a general model, which can be grounded on the behaviour of agents; the agents are the entities which populate the societies that are being regulated by legal systems and adopting them as the foundation for the representation is, in our view, the main contribution of this work. Of course, other agents do exist, most noticeably the judges and the members of the courts, that must decide about the behaviour of other agents; also judges (as all agents) are subject to norms, but they also have the role of assessing cases; this is a particular activity, that is not addressed in this paper, but that can be arguably assimilated to a kind of diagnostic reasoning: given some evidence (symptoms), classify the situation according to the categories appearing in the system. Since, in doing that, the court must respect some norms, it seems reasonable to assume that the model applies to it equally well, provided that a suitable representation is associated with the process of assessing the case.

Finally, the complexity of the representation. It is a matter of taste if a graphical representation (as the one we have adopted in this paper) is more or less readable than an equivalent representation respecting the linear syntax of a terminological language (see Appendix A). However, independently of the graphical way the ontology is presented, there is the usual contrast between the expressivity of the language and the computational complexity of the reasoning procedures. We have argued in favour of a language including some constructs that makes the language computationally hard. Although this is not the focus of the paper, we must stress that this appears to be unavoidable, unless one is able to show that the same concepts can be represented with simpler constructs, or that the knowledge expressed by these concepts can be left out of the ontology without loss of information. We believe that neither is the case, so that we hope that suitable heuristics can help a reasoner to obtain results within acceptable time limits.

Notes

¹ A complex network of physiological causes and effects with respect to “I feel bad”.

² “*Since the norm is the most salient construct of the legal domain, we have selected it as the point of departure for our intermediate representation*” (Kralingen et al. 1999, p. 1128).

³ In the text, we sometime use the category names in the plural. Although this is imprecise, we believe it makes the reading more fluent.

⁴ By ‘considering’, we mean here the process of moving around (e.g. from LTM to STM) pieces of descriptions.

- ⁵ More formally, by taking relations as functions, we have: $\forall x$ [Intention(x) \rightarrow (in-mind-of(x) = has-agent(p-effect-of(x)))] Of course, one can preserve relations, but the formula becomes less readable.
- ⁶ Often called *explicit permissions* (Boella and van der Torre 2003).
- ⁷ Note, however, that in Gangemi et al. (2003), Duty, Right, etc. are subcategories of Regulation, i.e., sisters of *Law* and *Obligation*.
- ⁸ For different positions on the role of sanctions in law, see Kelsen (1945) and Hart (1961).
- ⁹ The term *topic* is inspired to Ross (1968)'s "theme".
- ¹⁰ The fact that only DUTY appears in the figure is not intended to mean that all other relations are reduced to DUTY + polarity. Only DUTY is relevant in this example, but the other relations are there as sisters of DUTY.
- ¹¹ Since *Planning* is an *Action*, it is the domain of the *has-agent* relation; however, since *Planning* is a mental action, its agent could more properly be called *author* rather than *actor*, since the effect is a mental representation. Consequently, we could say that the *author* of a plan and its *actor* do coincide. The 'invented' plan could actually involve more than one 'actor', but we simplify the matter by calling *actor* the primary participant (the host of the dinner, and not the cook, the waiters, etc).
- ¹² We are dealing here with 'normative' rules. In other cases, the sanction may be replaced by a reward.
- ¹³ This is consistent with Hart (1961), p. 33, who characterizes power-conferring norms as recipes for creating duties.
- ¹⁴ Actually, an organization as a *Company* can also create an internal system of rules that constrain and organize the behaviour of its personnel. So, in principle, it can also be seen as an *Authority*. A partial solution is to create a special type of *LegalSubject* which refers to aggregates having a rich internal organization, some of whose internal organism have the role of *Authorities*.
- ¹⁵ As one referee pointed to us, the specification of what may constitute a *Good* from a legal point of view requires a more careful analysis. It is unclear, for instance, whether a *Description* may be a *Good*, while it is certain that some kinds of *Activities* (i.e., services) certainly are *Goods*. Since the focus of this example is on *Fruits*, and, at least in principle, any *Good* can produce *Fruits*, we leave open the problem of a better characterization of *Goods*.
- ¹⁶ On this point see also Bench-Capon and Visser (1996).
- ¹⁷ It is worth stating here again that we are using the world 'ontology' a-critically. Breuker pointed to us that "Planning is a problem solving task, so the perspective is an epistemological, rather than an ontological one". This comment reflects the view that any 'role' taken by a physical entity in a process concerns epistemology, i.e., the way people conceptualize the process. We share Breuker's view, but we cannot disregard that mental processes, as planning, have, as a physical substratum, the agents' brains, which are ontological objects. However, the literature in the field of philosophy of mind concerning the relationships between mind and brain is vast and growing at increasing speed. Consequently, we simply declare that we agree that a 'description' of mental processes in terms of intentionality, beliefs, etc. should more properly fall within the domain of epistemology, how-

ever, we keep the term ‘ontology’ as the only term that can be used without an in-depth discussion, which is outside the scope of this article.

Appendix A: Example of LOOM representation

In order to give an idea of the actual meaning of the structures presented in the figures, we report here a LOOM version of the contents of Figure 10, assuming that the semantics of LOOM is clearly defined (see [<http://www.isi.edu/isd/LOOM/LOOM-HOME.html>]). Some comments are reported within the definitions, using the: annotations slot. In order to simplify the match between the figure and the LOOM definition, we have split the latter into three boxes, which have no representational function. The fourth box reports an example of a LOOM query; when executed, it retrieves all civil fruits known to the system. This last box has been included to show that, although LOOM does not allow for role constraints, they can be enforced in a query, as constraints on the data to be retrieved.

PropertyRel, and its roles:

```
(defrelation source
  :constraints (and (:domain CivilFruit) (:range ExploitedGood))
  :annotations ('source' is a binary relation between a 'CivilFruit' and
    an 'ExploitedGood'))
(defrelation owner
  :constraints (and (:domain PropertyRel) (:range Legally-Constructed-Person))
  :annotations ('owner' is a binary relation between a 'PropertyRel' and
    a 'Legally-Constructed-Person'))
(defrelation ownee
  :constraints (and (:domain PropertyRel) (:range Asset))
  :annotations ('ownee' is a binary relation between a 'PropertyRel' and
    an 'Asset'))
(defrelation prop-of
  :is (inverse ownee)
  :annotations (the relation 'prop-of' is the inverse relation of
    'ownee': x,y (prop-of (x,y) ownee (y,x)))
)
(defconcept PropertyRel
  :is (and Relation
    (at-least 1 owner)
    (at-least 1 ownee))
  :annotations ('PropertyRel' is a subcategory of 'Relation'; for each
    instance of 'PropertyRel' there is at least one instance
    of 'owner' and at least one instance of 'ownee'))
```

Goods and Fruits:

```
(defconcept Good)
(defconcept Asset
  :is SubjectToProperty
  :annotations ('Asset' is a subcategory of 'SubjectToProperty'))
(defconcept CivilFruit
  :is (and Asset Fruit)
  :annotations (a 'CivilFruit' is both an 'Asset' and a 'Fruit';
    x (CivilFruit (x) Asset (x) Fruit (x)))
```

Contracts between two subjects:

```
(defrelation contractor
  :constraints (and (:domain Contract#2) (:range Legally-Constructed-Person))
  :annotations ('contractor' is a binary relation between a 'Contract#2'
    and a 'Legally-Constructed-Person'))
(defrelation giver
  :is contractor
  :annotations ('giver' is a subrelation of 'contractor'))
(defrelation receiver
  :is contractor
  :annotations ('receiver' is a subrelation of 'contractor'))
(defrelation exchanged
  :constraints (and (:domain Contract#2) (:range Asset))
  :annotations ('exchanged' is a binary relation between a 'Contract#2'
    and an 'Asset'))
(defrelation primary
  :is exchanged
  :annotations ('primary' is a subrelation of 'exchanged'))
(defrelation secondary
  :is exchanged
  :annotations ('secondary' is a subrelation of 'exchanged'))
(defconcept Contract#2
  :is (and JurPlan
    (exactly 2 contractor)
    (exactly 1 giver)
    (exactly 1 receiver)
    (exactly 2 exchanged)
    (exactly 1 primary)
    (exactly 1 secondary)
    (exactly 2 effect)
    (some effect PossessX)
    (some effect PropertyY))
  :annotations ('Contract#2' is a subcategory of 'JurPlan', with exactly
    two 'contractor', one 'giver', one 'receiver', two
    'exchanged', one 'primary', one 'secondary', and two
    'effects'; some effect is a 'PossessX', some
    is a 'PropertyY'))
(defconcept PropertyX
  :is (:the (:inverse effect) Contract#2))
(defconcept PossessY
  :is (:the (:inverse effect) Contract#2))
```

Retrieving data about fruits:

```
(defun get-frutti ()
  (retrieve (?c ?g ?r ?o ?p ?f ?n)
    ; *** retrieve all 7-tuples such that
    ; ?c is a 'Contract#2'
    ; ?g is its 'giver'
    ; ?r is its 'receiver'
    ; ?o is its 'primary'
    ; ?f is its 'secondary'
    ; ?n is the resulting 'PossessX'
    (and (Contract#2 ?c)
      (giver ?c ?g)
      (receiver ?c ?r)
      (primary ?c ?o)
      (PropertyRel ?p)
      (ownee ?p ?o)
      (owner ?p ?g)
      (secondary ?c ?f)
      (PossessX ?n)
      ((:inverse effect) ?n ?c)
      (possessor ?n ?g)
      (possessee ?n ?f) (CivilFruit ?f)
      (source ?f ?o)))
```

Appendix B: The planning process

In this appendix, we present a simple model of the planning process. This has to be viewed as a preliminary attempt at an ontological representation of a complex activity as planning. We describe here the features that are depicted in Figure 12. Note that the content of the figure is a sort of procedure; although a procedure is inherently a ‘procedural’ representation of a process, its expression in a programming language also has a declarative (descriptive) flavour, since it can be inspected as a static object (for instance by a compiler). In order to provide an ontological representation, we need the equivalent of some basic programming concepts, as input parameter,

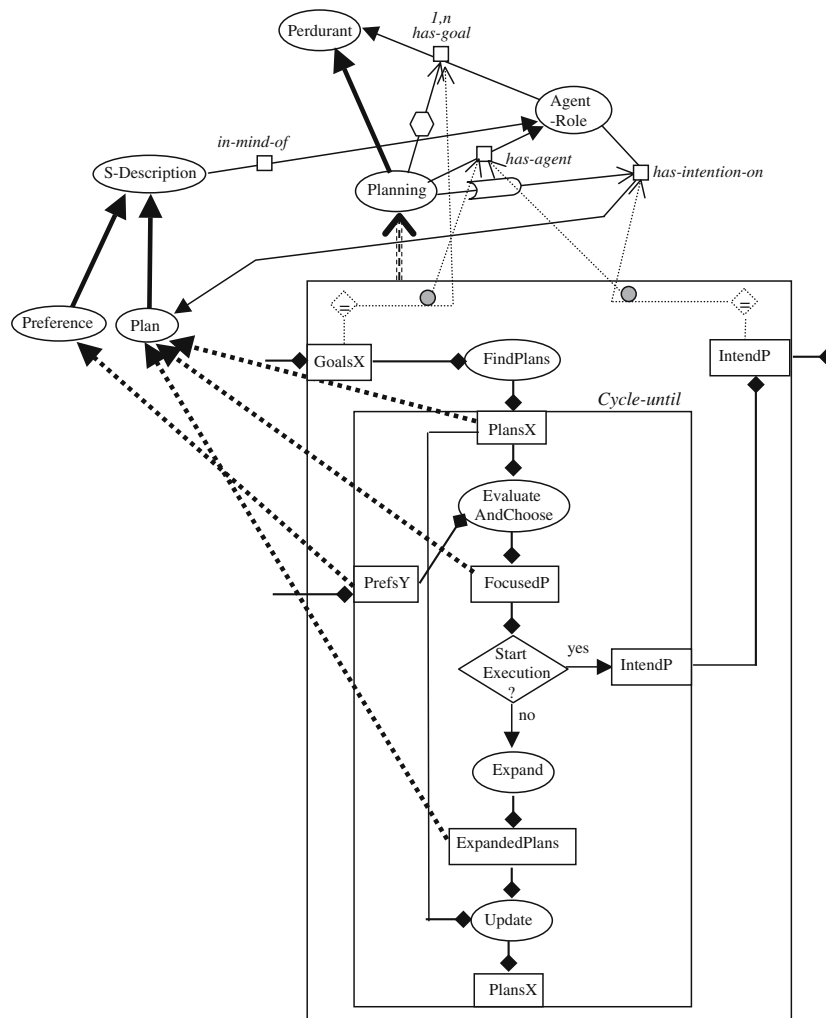


Figure 12. Structure of the Planning activity.

variable, sequence of operations, loop, etc. A first proposal appears in DOLCE in the section related to the *Task* concept (Masolo et al. 2003, pp. 238–241). Here, we have adopted the following graphical conventions:

1. *Planning* is linked to its ‘body’ (the large box) via a dashed arrow. This notion corresponds to the ‘decomposition’ of a complex action in its steps. In DOLCE, we have that a *Course* (which is a *Description*) *sequences* a *Perdurant*. Since a *Task* is a subcategory of *Course*, the large box in the figure can be taken as the complex *Task* that *sequences* (the dotted upward large arrow) the *Planning Perdurant*.
2. The sequence of operations is represented as lines ending in a black diamond. Some lines having this form enter the ‘body’ box: they are the input data, and the small box they enter is the parameter name (ex. *GoalsX*); some of these arrows exit the box: they are the result (the return values).
3. The sequence of steps is an alternation of ovals (e.g. *FindPlans*) and rectangles (e.g. *PlansX*); the former are sub-procedures (steps, or *Sub-Tasks*), while the latter are intermediate results (local variables).
4. In particular, *FindPlans* retrieves from the knowledge of the planning agent all the *Plans* that have as *p-effects* one or more of the *Goals* in *GoalsX*.
5. The *Cycle-Until* box expresses a loop (in DOLCE, it is a subcategory of *Complex-Task*): all the steps inside the box are repeated until the exit condition (*StartExecution?*) is satisfied. The predicate is assumed to be true if a single plan has remained in the focus (*FocusedP*) and the initial steps of the focused plan are executable; this last condition holds if that step refers to a ‘basic action’, i.e., to an action that can drive the actuators of the agent.
6. *EvaluateAndChoose* aims at evaluating the foreseeable utility of the different alternatives (see Section 2.4); it gets as input the preferences of the planning agent (*PrefsX*).
7. *Expand* takes into account all remaining focused plans (which may be more than one in case no commitment has been made), and finds ways for executing non-basic actions. After *Expand*, the resulting *Plans* are more specific than the original ones.
8. It is specified (via the usual constraints on roles) that the input goals (*GoalsX*) are the ones of the planning agent. A similar specification should be given for the preferences; we have omitted it because of space constraints.
9. Finally, it is specified that the final plan which is the result of the whole process (*IntendP*) is the one that becomes the intention of the planning agent (*has-intention-on*). This is the connection (effects of planning) to the concepts appearing in Figure 4.

References

- Alchourron, C. (1993). Philosophical Foundations of Deontic Logic and the Logic of Defeasible Conditionals. In Meyer, J.-J. and Wieringa, R. (eds.), *Deontic Logic in Computer Science: Normative System Specification*, 43–84. John Wiley & Sons.
- Allen, L. E. and Saxon, C. S. (1995). Better Language, Better Thought, Better Communication: The A-Hohfeld Language for Legal Analysis. In *Proceedings of 5th International Conference on Artificial Intelligence and Law*. ACM Press.
- Bench-Capon, T. J. M. and Visser, P. R. S. (1996). Deep Models, Ontologies, and Legal Knowledge Based Systems. In van Kralingen, R. W., van den Herik, H. J., Prins, J. E. J., Sergot, M. and Zeleznikow, J. (eds.), *Proceedings of JURIX '96: Foundations of Legal Knowledge Systems*, 3–13. Tilburg University Press.
- Bench-Capon, T. J. M. (2001). Task Neutral Ontologies, Common Sense Ontologies and Legal Information Systems. In *Proceedings 2nd Workshop on Legal Ontologies at Jurix*, 15–19.
- Boella, G., Lesmo, L., and Favali, L. (2001). The Definition of Legal Relations in a BDI Multiagent Framework. In *Proceedings of AI*IA 01*, 225–236. Berlin: Bari, Springer Verlag.
- Boella, G. and Lesmo, L. (2002). A Game Theoretic Approach to Norms. *Cognitive Science Quarterly* 2: 492–512.
- Boella, G. and van der Torre, L. (2003). Obligations and Permissions as Mental Entities. In *Proceedings of IJCAI Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions*. Acapulco.
- Boer, A., Hoekstra, R., and Winkels, R. (2001). The CLIME Ontology. In *Proceedings of 2nd Workshop on Legal Ontologies at Jurix*, 37–47.
- Brachman, R. J. and Schmolze, J. (1985). An Overview of kl-one. *Knowledge Representation System*. *Cognitive Science* 9: 171–216.
- Breuker, J. (1997). André Valente and Radboud Winkels: Legal Ontologies: A Functional View. In *Proceedings of the First International Workshop on Legal Ontologies*, 23–36.
- Breuker, J. and Winkels, R. (In press). Use and Reuse of Legal Ontologies in Knowledge Engineering and Information Management. *AI & Law*.
- Carbonell, J. G., Etzioni, O., Gil, Y., Joseph, R., Knoblock, C. A., Minton, S., and Veloso, M.M. (1991). PRODIGY: An Integrated Architecture for Planning and Learning 2: 51–55.
- Gangemi, A., Pisanelli, D. M., and Steve, G. (2001). A Formal Ontology Framework to Represent Norm Dynamics. In *2nd International Workshop on Legal Ontologies*. University of Amsterdam.
- Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., and Schneider, L. (2002). Sweetening Ontologies with DOLCE. In *Proceedings of EKAW 2002*. Spain: Siguenza.
- Gangemi, A., Sagri, M. T., and Tiscornia, D. (2003). Metadata for Content Description in Legal Information. In *Proceedings of LegOnt 2003 Workshop*. Edinburgh.
- Gil, Y. and Blythe, J. (2000). PLANET: A Shareable and Reusable Ontology for Representing Plans. In *Proceedings of AAAI 2000 Workshop on Representational Issues for Real-world Planning Systems*.
- Guarino, N. (1998). Some Ontological Principles for Designing Upper Level Lexical Resources. In *Proceedings of First International Conference on Language Resources and Evaluation (LREC 1998)*. Granada.
- Guarino, N. and Welty, C. (2000). A Formal Ontology of Properties. In *Proceedings of EKAW 2000*. 97–112.
- Guarino, N. and Welty, C. (2002). Evaluating Ontological Decisions with OntoClean. *Communications of the ACM* 45: 61–65.

- Hage, J. and Verheij, B. (1999). The Law as a Dynamic Interconnected System of States of Affairs: a Legal Top Ontology. *International Journal of Human-Computer Studies* 54: 1043–1077.
- Hart, H. (1961). *The Concept of Law*. Clarendon Press: Oxford.
- Herrestad, H. and Krogh, C. (1995). Obligations Directed from Bearers to Counterparties. In *Proceedings of 5th International Conference on Artificial Intelligence and Law*. ACM Press.
- Hohfeld, W. N. (1919). *Fundamental Legal Conceptions as Applied in Judicial Reasoning and Other Legal Essays*. New Haven, Conn.: Yale University Press.
- Kelsen, H. (1945). *General Theory of Law and State* (Wedberg trans.). Russell & Russell: New York.
- van Kralingen, R. W., Visser, P. R. S., Bench-Capon, T. J. R., and van den Herik, H. J. (1999). A Principled Approach to Developing Legal Knowledge Systems. *Int. J. Human-Computer Studies* 54: 1127–1154.
- Lee, J., Gruninger, M., Jin, Y., Malone, T., Tate, A., and Yost, G. (1998). Process Interchange Format and Framework. In Uschold, M. and Tate, A. (eds.), *The Knowledge Engineering Review 13. Special Issue on Ontologies*, <http://ccs.mit.edu/pif/>.
- Lehmann, J., Breuker, J., and Brouwer, B. (2003). Causation in AI & Law. In *ICAIL 2003 Workshop on Legal Ontologies & Web Based Legal Information Management*, 1–34. Edinburgh, June 2003.
- Masolo, C., Borgo, S., Gangemi, A., Guarino, N., and Oltramari, A. (2003). *WonderWeb Deliverable D18. Ontology Library. Final Report* (ver. 1.0, 31-12-2003).
- MacGregor, R. M. (1991). Using a Description Classifier to Enhance Deductive Inference. In *Proceedings of 7th IEEE Conference on AI Applications*, 141–147.
- MacGregor, R. M. and Brill, D. (1992). Recognition Algorithms for the Loom Classifier. In *Proceedings of 10th National Conference on Artificial Intelligence (AAAI 92)*, 774–779.
- Mommers, L. (2001). A Knowledge-based ontology of the legal domain. *Proceedings of the Second Workshop on Legal Ontologies at Jurix 2001*: 1–7.
- Pollack, M. E. (1990). Plans as Complex Mental Attitudes. In Cohen, P. R., Morgan, J. and Pollack, M. E. (eds.), *Intentions in Communication*, 77–103. MIT Press.
- Polyak, S. and Tate, A. (1998). *Planning Initiative: Shared Planning and Activity Representation – SPAR, Version 0.1a*. DARPA Internal Report, April 1998.
- Polyak, S. and Tate, A. (1999). *Planning Initiative: Shared Planning and Activity Representation – SPAR*. DARPA Internal Report, March 1999.
- Prakken, H. and Sergot, M. (1997). Dyadic Deontic Logic and Contrary-to-duty Obligations. In Nute, D. (ed.), *Defeasible Deontic Logic*, 223–262. Kluwer.
- Ross, A. (1968). *Directives and Norms*. Humanities Press: New York.
- Sen, A. (1999). *Development as Freedom*. Oxford University Press.
- Smith, S. F. and Becker, M. A. (1997). An Ontology for Constructing Scheduling Systems. In *Proceedings of AAAI-97 Spring Symposium on Ontological Engineering*.
- Stuckenschmidt, H., Stubkjaer, E., and Schlieder, C. (2001). Modeling Land Transactions: Legal Ontologies in Context. In *2nd Workshop on Legal Ontologies at Jurix 2001*, 58–66.
- Tiscornia, D. (2001). *Ontology-driven Access to Legal Information*. In *DEXA 2001 – Workshop LISA, Legal Information Systems and Application*. Munich, September 2001.
- Valente, A., Breuker, J., and Brouwer, B. (1999). Legal Modelling and Automated Reasoning with ON-LINE. *Int. J. Human-Computer Studies* 54: 1079–1125.
- Visser, P. R. S. and Bench-Capon, T. J. M. (1998). A Comparison of Four Ontologies for the Design of Legal Knowledge Systems. *Artificial Intelligence and Law* 6: 27–57.

- van der Torre, L. and Tan, Y.-H. (1995). Cancelling and Overshadowing: Two Types of Defeasibility in Defeasible Deontic Logic. In Proceedings of 14th International Joint Conference on Artificial Intelligence (IJCAI' 95), Morgan Kaufman, 1525–1532.
- von Wright, G. H. (1951). Deontic Logic. *Mind* 60: 1–15.
- Welty, C. and Guarino, N. (2001). Supporting Ontological Analysis of Taxonomic Relationships. *Data and Knowledge Engineering* 39: 51–74.
- Winkels, R., Bosscher, D. J. B., Boer, A., and Breuker, J. (1999). Generating Exception Structures for Legal Information Serving. In Proceedings of ICAIL 99. New York, 182–189.
- Winkels, R., Bosscher, D. J. B., Boer, A., and Hoekstra, R. (2000). Extended Conceptual Retrieval. In Breuker, J. A. et al. (eds.), *Legal Knowledge and Information Systems 2000*, 85–98. Amsterdam: IOS Press.
- Zelevnikov, J. and Stranieri, A. (2001). The use of Legal Decision Support Systems at Victoria Legal Aid. In Proceedings of ISDSS2001, Sixth International Conference on Decision Support Systems, 186–192. London: Brunel University.