Chapter 16
Norms and Argumentation

Nir Oren, Antonino Rotolo, Leendert van der Torre, and Serena Villata

16.1 Introduction

The study of norms and argument has become increasingly connected in recent times, particularly in domains such as law, knowledge representation, ethics, linguistics and, most recently, in various agreement technologies (see also Part V in this book). Here, norms are used to set the space of legal agreements (or commitments) and argumentation is used to choose among the possible agreements (Billhardt et al. 2011). Moreover, we may consider norms setting not only the scope of possible legal agreements, but also the way we can choose among these possible agreements.

In law, Bench-Capon et al. (2010) present how argumentation theory has been used in legal reasoning. For instance, legal disputes arise out of a disagreement between two parties and may be resolved by presenting arguments in favor of each party's position. These arguments are proposed to a judging entity, who will justify the choice of the arguments he accepts with an argument of his own, with the aim of convincing the public. The common conclusion shared by such works is that argumentation has the potential to become a useful "tool" for people working in the legal field. While legal practitioners typically believe that argumentation theory can

N. Oren (✉)
University of Aberdeen, Aberdeen, UK
e-mail: n.oren@abdn.ac.uk

A. Rotolo
CIRSFID – Faculty of Law, University of Bologna, Bologna, Italy
e-mail: antonino.rotolo@unibo.it

L. van der Torre
ICR Group, University of Luxembourg, Walferdange, Luxembourg
e-mail: leon.vandertorre@uni.lu

S. Villata
INRIA Sophia Antipolis, Sophia Antipolis, France
e-mail: serena.villata@inria.fr

S. Ossowski (ed.), *Agreement Technologies*, Law, Governance and Technology Series 8, DOI 10.1007/978-94-007-5583-3_16,

simply be used to deduce consequences from a set of facts and legal rules, or to detect conflicts within such sets, argumentation can offer much more. Following the example proposed by Bench-Capon et al. (2010), a case is not a mere set of facts, but can be seen as a story told by a client to their lawyer. The first thing the lawyer does is to interpret this story in a particular legal context. The lawyer can interpret the story in several different ways, and each interpretation will require further facts to be obtained. Then the lawyer has to select one of the possible interpretations, must provide arguments to persuade the judging entity of the client's position, and has to rebut any further objection. The major topics that emerge as relevant in norms and argumentation include, among others, case based reasoning (Ashley 1990; Rissland et al. 1993), arguing about conflicts and defeasibility in rule based systems (Prakken 1993; Prakken and Sartor 1996; Sergot et al. 1986), dialogues and dialectics (Gordon 1993), argument schemes (Bex et al. 2003; Gordon and Walton 2009), and arguing about the success of attacks (Farley and Freeman 1995; Prakken et al. 2005).

In this chapter, we highlight the future challenges in the research area of norms and argumentation. We show that existing work on norms and argumentation, some of which has been identified above, can be categorized into two different classes, namely (i) arguing about norms, and (ii) norms about argumentation. The former includes the greater part of existing work in the area of norms and argumentation, such as approaches which aim at resolving conflicts and dilemmas (in particular examining how norms interact with other norms), arguing about norm interpretation and dynamics, arguing about norm adoption, acceptance and generation, representing norm negotiation, and arguing about contracts. In spite of all the literature on these topics, several challenges still remain to be addressed and resolved. For instance, new frameworks where individuals can discuss the merits and effects of the norms to be adopted by a society, or the introduction of richer preference models to detect and reason about norm interactions still provide fertile ground for additional research. At the moment far less work exists dealing with norms about argumentation. This topic aims to address the challenges of dialogue and debate protocols, reasoning about epistemic norms, and enforcement models of the burden of proof. In this category of work, open questions remain regarding the introduction of new techniques to verify whether a virtual agent complies with an epistemic norm, and the development of tools able to support the judging entities and the lawyers in the enforcement of burden of proof. Finally, besides the norms about argumentation and arguing about norms, direct formal relations between deontic logic—in particular input/output logic—and abstract argumentation have been considered (Bochman 2003, 2005), leading to a number of additional challenges.

16.2 Arguing About Norms

In order to determine how argumentation can be useful when dealing with norms, we must first examine the concept of a norm in more detail as done in Part III. Searle (1997) distinguished between two types of norms, referred to as regulative

and constitutive norms (Boella and Van der Torre 2004b). Such norms form a group or society's *normative system*, and directly or indirectly constrain the behaviour of the society. Regulative norms modify an individual society member's behaviour by *obliging, permitting* or *prohibiting* certain states of affairs from occurring. If such norms are *violated*, additional regulatory norms, referred to as *contrary-to-duty* norms, can come into effect, leading to sanctions against the violators.[1] Constitutive norms describe the society through a mapping from *brute facts* to societal or institutional facts via a *counts-as* relation (cf. the societal concept of marriage); by identifying societal power structures through notions such as roles, and by specifying how the normative system itself can be modified (e.g. by allowing an entity taking on some role to modify some of the norms).

While regulative and constitutive norms affect a society in very different ways, both types of norms must be recognized by the members of a society in order to affect the society. The question then immediately arises as to how norms can be created in such a way so as to be recognized by members of a society. One possibility involves the assumption of some underlying system of sanctions, rewards and existing norms that allow entities taking on the role of a legislator to simply insert new norms into the society, with these norms recognized by all members of a society (Artikis et al. 2009; Gelati et al. 2004; Oren et al. 2010). The injection of norms into a system by a system designer can be seen to fall into this category of norm creation. Another alternative, espoused by work such as (Boella and van der Torre 2004a) takes a game theoretic view of *norm emergence*; norms come into being and are accepted by the society when adhering to the norm results in a Nash equilibrium for the members of the society. The computational overheads of computing a Nash equilibrium can make the latter approach infeasible, while the former (effectively) assumes some dictatorial power in the system. An approach in which individuals within the society can debate the merits of a norm, discuss its effects, and persuade others as to the utility of its adoption could provide a level of societal modelling and control not found in existing systems. Such *argumentation based norm creation* could be used as both a modelling tool (e.g. to model how laws are generated), and as a technique to reason about the effectiveness of proposed norms.

Having briefly discussed the role of arguments in norm creation and recognition, we will now consider the advantages that an argumentation based approach can bring to various aspects of both constitutive and regulative norms.

16.2.1 Argumentation and Constitutive Norms

Constitutive norms capture two aspects of a society, namely an ontological aspect mapping brute facts (e.g. the recital of vows) to societal facts (e.g. a couple being viewed as married according to the society), and the power structures found in the

[1]This view of regulative norms encapsulates social norms, which impose requirements on behaviour with no explicit sanction.

society. Having discussed the latter aspect in the previous section, we now examine the contributions argumentation theory can make to the former.

While members of a society should agree on the definition of societal facts (i.e. share an ontology), this is not always the case. Even in human societies, legal cases often rest on the judge's interpretation of some definition. The question then arises as to how the ontology representing an individual's view of the societal facts can be aligned with those of all other member of the society. Most existing approaches to attacking this *ontology alignment problem* originate from the Semantic Web community, and are neither able to deal with conflicting perspectives on the meaning of terms, nor cater for non-Description Logic based knowledge representation schemes. Several argumentation based approaches have been proposed to tackle ontology alignment and matching. For example, dos Santos et al. (2009) represents the outcomes of several matching tools within an abstract argument framework; the extensions of this framework are then used to define the final ontology alignment process. Laera et al. (2007) takes a semantic approach, with agents exchanging arguments over the properties of the ontology in order to agree on a compromise. Evaluations of these argument based approaches have shown that they perform as well as the best standard ontology mapping techniques.

While promising, such techniques must overcome important restrictions before they are suitable for the alignment of constitutive norms. These include the assumption that only two agents are attempting to perform ontology alignment, and the requirement for a central ontology mapping repository.

16.2.2 Argumentation and Regulative Norms

Given some conditions, regulative norms impose obligations, permissions and prohibitions on members of a society. An individual must be able to recognise which norms should, according to the society, affect its behaviour. Since a single society can have multiple norm creators, there is no guarantee that the norms affecting the individual are consistent, and it must therefore be possible to detect *normative conflicts*. Now, one approach to dealing with such normative conflict involves ignoring all but one of the conflicting norms. However, an individual could ignore a norm for other reasons. For example, violating a norm could allow the individual to achieve an important goal. Therefore, rather than treating norms in isolation, normative reasoning must form part of practical reasoning.

The violation of a regulative norm typically carries with it the threat of a sanction. In order to make good on such a threat, a society must be able to monitor the compliance of an individual with regards to the norms affecting it. Argumentation has a role to play in addressing each of the aspects of regulative norms discussed so far. In the remainder of this section, we will examine each of these aspects in more detail, describe existing work dealing with the aspect, and identify how argumentation could be used to address remaining challenges.

16.2.2.1 Identifying Normative Constraints

Regulative norms most commonly impose conditional obligations on an individual or set of individuals. Such conditional obligations identify what state of affairs *should* hold when some situation occurs. Norms can also cause prohibitions to come into force. These are similar to obligations, but identify the state of affairs that *should not* hold. Finally, regulative norms can also instantiate permissions. Makinson and van der Torre (2003) identifies three separate types of permissions, namely weak, strong and dynamic permissions. Weak permissions—stating that what is obliged is permitted—exist implicitly within a normative system, and are not further discussed here. Strong permissions identify states of affairs which cannot be prohibited, while dynamic permissions are used to *derogate* obligations and prohibitions. The inclusion of the latter type of permission into a normative system therefore results in a defeasible system (Horty 1997); abstract argument frameworks can be used to represent such systems, with norms within an extension representing constraints upon the individual, and norms outside the extension being derogated.[2] Instantiations of such an abstract framework via a defeasible logic, in the vein of Governatori and Rotolo (2010), allows for complex normative reasoning to take place, incorporating concepts such as deadlines, norm violation, and norm fulfilment. The main direction of future work in this area involves further refining existing frameworks. Another possible area of investigation would make use of argument schemes to reason about why certain norms are, or are not in force, based on uncertain evidence from the domain. Preliminary work in this direction was discussed in Oren et al. (2008).

16.2.2.2 Detecting and Dealing with Normative Conflict

If complying with any norm will result in some other norm being violated, then the set of norms is in *normative conflict*. While some forms of normative conflict are easy to detect (e.g. an obligation to achieve a, and a prohibition on achieving a), others require explicit domain knowledge (such as mutual exclusivity between actions). Now permissions may *derogate* prohibitions and obligations, temporarily alleviating the conflict. Work such as (Governatori and Rotolo 2008) makes use of a defeasible framework to reason about such derogations.

Argumentation is designed to represent and reason about conflicts, for instance in Oren et al. (2008) an abstract argument based approach selecting which norms to violate in the presence of normative conflict was described. Possible enhancements to this argumentation based approach include richer preference models, as well as logics for reasoning about norm interactions in order to reason about, and detect, future potential norm conflicts.

[2] If such an abstract framework contains multiple extensions, then the potential exists for normative conflict to arise, as discussed later.

We have described how argumentation can be used to reason about normative conflict, allowing an agent to decide how to act in its presence. This reasoning process can be generalised further, namely by reasoning about how to act in the presence of conflicts between an agent's norms and its goals.

The BOID architecture (Broersen et al. 2001) utilised defeasible rules and a priority relation in order to perform practical reasoning in the presence of conflicts between obligations, intentions and desires. Work such as (Boella and van der Torre 2003) then investigated how argumentation could be used as the underlying reasoning mechanism within a BOID architecture. Advantages of argumentation in this context include the ability to include permissions as undercutting attacks, and naturally have agents influence each other's behaviour. Following this strand of work, Modgil and Luck (2009) makes use of extended argument frameworks to represent an agent's reasoning process when performing practical reasoning in the presence of conflicting desires and norms.

There are strong analogies between the practical reasoning problem and automated planning, and recent work (e.g. Toniolo et al. 2011) has proposed an argumentation based approach to planning in the presence of norms. Apart from its ability to handle conflict, the use of argumentation for practical reasoning enables easy integration of domain specific knowledge and inference (via argument schemes). The use of argumentation in this context offers the possibility of improved computational performance, potentially allowing for new reasoning heuristics. Preliminary work such as Medellin-Gasque et al. (2011) is now underway on encoding such domain specific argument schemes, and is possibly the most exciting area for future research.

16.2.2.3 Monitoring Norms

Until now, we have examined how an agent can reason about its norms in order to decide which norms should affect its behaviour, and how to act in the presence of these norms. Now if an agent violates a norm, then it can be sanctioned by other agents within the society. To this end, some mechanism is required to monitor the norm compliance (or lack thereof) of agents. Several such mechanisms, (e.g. DIO(DE)^2 van der Torre and Tan 1999) have been proposed to perform such monitoring. One important requirement for such a monitoring system is the ability to handle both conflicting and uncertain evidence; the former because agents with a vested interest in sanctioning the monitored agent could attempt to deceive the monitor, and the latter due to a lack of omniscience on the part of the monitor. This problem of norm monitoring is analogous to the problem of contract monitoring (Daskalopulu et al. 2002). In Oren et al. (2007, 2008), argumentation based approaches to reasoning about uncertain evidence were proposed. More generally, a large body of work exists regarding arguing in the presence of uncertainty (e.g. Haenni et al. 2001 and McBurney and Parsons 2000), and challenges here include identifying argument schemes which reason about uncertainty, how to weigh up conflicting uncertain evidence and so on.

16.2.2.4 Norm Dynamics

Obligations can change while the normative system remains the same. For example, due to change in the world or in the agents' knowledge and beliefs, new obligations can be detached from the norms, or an agent can delegate one of its obligations to another agent. This change of obligations and permissions over time is a relatively clear and well studied subject, investigated mostly in the 1970s and 1980s. Moreover, a code of regulations is itself not static either, but changes over time. For example, a legislative body may want to introduce new norms or to eliminate some existing ones. To study how norm change is different from how obligation change, and how these two are related, we have to address topics such as:

- Norm revision and contraction, e.g. change of legal code;
- Norm evolution, e.g. change of social norms;
- Merging normative systems, e.g. the merge of companies.

Note that we presuppose a distinction between norms and obligations, which is too often ignored. Norms, imperatives, promises, legal statutes, and moral standards are usually not viewed as being true or false. For example: "John, leave the room!" and "Mary, you may enter now" do not describe, but demand or allow a behaviour on the part of John and Mary. Lacking truth values, norms cannot be premise or conclusion in an inference, be termed consistent or contradictory, or be compounded by truth-functional operators. The usual way out is to say that "John is obliged to leave the room" describes the obligation which follows from the prescriptive "John, leave the room!" Makinson (1998) raises the question: How can deontic logic be reconstructed in accord with the philosophical position that norms are neither true nor false?

The derived problem is: How to formalize the relation between norm change and obligation change?

Little work exists on the logic of the revision of a set of norms. To the best of our knowledge, Alchourrón and Makinson were the first to study the changes of a legal code (Alchourrón and Makinson 1982). The addition of a new norm n causes an enlargement of the code, consisting of the new norm plus all the regulations that can be derived from n. Alchourrón and Makinson distinguish two other types of change. When the new norm is incoherent with the existing ones, we have an *amendment* of the code: in order to coherently add the new regulation, we need to reject those norms that conflict with n. Finally, *derogation* is the elimination of a norm n together with whatever part of G implies n.

Some of the AGM (Alchourrón et al. 1985) axioms seem to be rational requirements in a legal context, whereas they have been criticized when imposed on belief change operators. An example is the *success* postulate, requiring that a new input must always be accepted in the belief set. It is reasonable to impose such a requirement when we wish to enforce a new norm or obligation. However, it gives rise to irrational behaviors when imposed to a belief set, as observed for instance in Gabbay et al. (2003).

We now want to turn to another type of change, that is the aggregation of regulations. This problem has only recently been addressed in the literature and therefore the findings are still very partial. The aggregation of regulations can be addressed using argumentation techniques developed for merging argumentation frameworks.

The first noticeable thing is the lack of general agreement about where the norms that are to be aggregated originate. Some works focus on the merging of conflicting norms that belong to the same normative system, while other works assume that the regulations to be fused belong to different systems. The first situation seems to be more a matter of coherence of the whole system rather than a genuine problem of fusion of norms. However, such approaches have the merit of revealing the tight connections between fusion of norms, non-monotonic logics and defeasible deontic reasoning.

We have seen that the initial motivation for the study of belief revision was the ambition to model the revision of a set of regulations. On the contrary, the generalization of belief revision to *belief merging* is exclusively dictated by the goal to tackle the problem—arising in computer science—of combining information from different sources. The pieces of information are represented in a formal language and the aim is to merge them in an (ideally) unique knowledge base. Can the belief merging framework deal with the problem of merging sets of norms?

The AGM framework has the advantage of being very abstract but works with theories consisting of simple logical assertions. For this reason, it is perhaps suitable to capture the dynamics of obligations and permissions, not of legal norms. In fact, it is essential to distinguish norms from obligations and permissions (Boella et al. 2009; Governatori and Rotolo 2010): the latter ones are just possible effects of the application of norms and their dynamics do not necessarily require to remove or revise norms, but correspond in most cases to instances of the notion of *norm defeasibility* (Governatori and Rotolo 2010). Very recently, some research has been carried out to reframe AGM ideas within rule-based logical systems, which take this distinction into account (Stolpe 2010). However, also these attempts suffer from some drawbacks, as they fail to handle the following aspects of legal norm change:

1. The law usually regulate its own changes by setting specific norms whose peculiar objective is to change the system by stating what and how other existing norms should be modified;
2. Since legal modifications are derived from these peculiar norms, they can be in conflict and so are defeasible;
3. Legal norms are qualified by temporal properties, such as the time when the norm comes into existence and belongs to the legal system, the time when the norm is in force, the time when the norm produces legal effects, and the time when the normative effects hold.

Hence, legal dynamics can be hardly modeled without considering defeasibility and temporal reasoning. Some recent works (see, e.g., Governatori and Rotolo 2010) have attempted to address these research issues. All norms are qualified by the above

mentioned different temporal parameters and the modifying norms are represented as defeasible meta-rules, i.e., rules where the conclusions are temporalized rules.

In this section we described how argumentation can be used with regards to several aspects of normative reasoning. Much of the future work in this area lies in making use of argumentation to extract, and utilize, the reasoning mechanisms specific to the normative domain. In the next section, we examine the dual of this approach. Namely, we examine how norms can be used to guide the process of argumentation.

16.3 Norms About Argumentation

In this section, we present and discuss some issues like debate protocols, and burden of proof where norms have the role of regulations on the argumentation process itself.

16.3.1 Dialogue and Debate Protocols

Norms about argumentation constitute what argumentation conceptually is and what it factually should be. Hence, such norms are supposed to provide a framework where the exchange of opinions makes sense, it rigorously takes place, looks acceptable and rational. This is done in argumentation theory by identifying, as mentioned in Chap. 13, formal requirements for at least three different layers: the logical layer, the dialectical layer, and a procedural layer (Prakken and Sartor 2002; Prakken and Vreeswijk 2002).

16.3.1.1 Logical Layer

The logical layer deals with the underlying language that is used to build arguments. Many languages and reasoning methods can be used for this purpose, such as deduction, induction, abduction, analogy, and case-based reasoning. If the underlying language refers to logic **L**, arguments can roughly correspond to proofs in **L** (Prakken and Sartor 2002). It may be argued that most argumentation systems are based on a *monotonic* consequence relation, since each single argument cannot be revised but can only be invalidated by *other* arguments (or better, counter-arguments) (Prakken and Vreeswijk 2002): it is the exchange of arguments and counter-arguments that makes the system non-monotonic. However, this is not strictly required: when the underlying logic is itself non-monotonic, an argumentation system can be simply seen as an alternative way to compute conclusions in that non-monotonic logic (Governatori et al. 2004).

Suppose we resort to a rule-based logical system where rules have the form $\phi_1, \ldots, \phi_n \Rightarrow \phi$ and represent defeasible norms. An argument for a normative conclusion ϕ can typically have a tree-structure, where nodes correspond to literals and arcs correspond to the rules used to obtain these literals; hence, the root corresponds to ϕ, the leaf nodes to the primitive premisses, and for every node corresponding to any literal ψ, if its children are $\psi_1, \ldots, \psi_n$, then there is a rule whose antecedents are these literals (Governatori et al. 2004).

Argumentation systems, however, do not need in general to specify the internal structure of their arguments (Dung 1995). In this perspective, any argumentation system $\mathscr{A}$ is a structure $(A, \rightsquigarrow)$, where A is a non-empty set of arguments and $\rightsquigarrow$ is binary attack relation on A: for any pair or arguments a and b in A, $a \rightsquigarrow b$ means that a attacks b. This leads us to discuss the dialectical layer.

16.3.1.2 Dialectical Layer

The dialectical layer addresses many interesting issues, such as when arguments conflict, how they can be compared and what arguments and conclusions can be justified.

Different types of attacks and defeat relations can apply to arguments. Pollock's (1995) original distinction between *rebutting* and *undercutting* is almost universally accepted in the argumentation literature (Prakken and Sartor 2002, 2004). An argument A_1 rebuts an argument A_2 when the conclusion of A_1 is equivalent to the negation of the conclusion of A_2. The rebutting relation is symmetric. For example, if arguments are built using rules representing norms (regulating, for example, smoking in public spaces), a conflict of this type at least corresponds to a clash between the conclusions obtained from two norms (for example, one prohibiting and another permitting smoking). The undercutting is when an argument challenges a rule of inference of another argument. This attack relation is not symmetric and occurs when an argument A_1 supporting the conclusion ϕ has some ground ψ but another argument A_2 states that ψ is not a proper ground for ϕ. To put it simply, if one builds an argument A_1 for ϕ using the rules $\Rightarrow \psi$ and $\psi \Rightarrow \phi$ but we contend that ψ is the case, then we undercut A_1.

16.3.1.3 Procedural Layer

The procedural layer considers the ways through which conclusions are dynamically reached by exchanging arguments between two or more players. Legal disputes, for instance, can be reconstructed in the form of dialogues, namely of players' dialectical moves (Gordon 1995; Prakken 2001). Disputes in turn are regulated by procedural rules stating what dialogue moves (claiming, challenging, conceding, etc.) are possible, when they are "legal", what effects the players get from them, and under what conditions a dispute terminates (Gordon 1995) (in general, see Walton and Krabbe 1995).

In this sense, the procedural layer is the place where the dimension of norms about argumentation and the one of argumentation about norms are strongly linked: complying with the norms that regulate, for example, deliberative processes or legal disputes is a guarantee of fairness and justice for those processes and disputes, but fairness and justice are supposed to be behind any legal, moral and social norms.

In general, different frameworks for the procedural layer can typically be thought in terms of defining different dialogue systems. In this perspective, Walton and Krabbe (1995) identified for example the following fundamental dialogue types: persuasion dialogue, negotiation, information seeking dialogue, deliberation, inquiry, and quarrel. Each of these types correspond to a different way through which the argumentation can dynamically take place.

An established method for characterizing dialogues specifies and regulates their dynamic development by the so called protocol, effect rules, and outcome rules (Barwise and Moss 1996; Prakken 2005). The protocol states what moves are allowed (the "legal" moves) at each point in a dialogue (turntaking and termination); the effect rules define, for each utterance, the consequences for the commitments of the players; the outcome rules state the outcome of a dialogue (for instance: in persuasion dialogues the outcome is establishing the winner and the loser, while in negotiations the outcome is some allocation of resources among the players).

A basic and fundamental question of the procedural layer regards how to govern and allocate the burden of proof (Prakken 2001). For example, basic legal dialogue protocols of 2-player in civil disputes are defined on account of the requirement that the plaintiff begins the dispute with his claim and has to propose, to win, at least one justified argument which supports such a claim. The burden of the defendant is not in principle the same, as it may be sufficient in most cases for her to oppose to the plaintiff argument moves that are only defensible counter-arguments. The concept of legal burden of proof is very complex and its logical treatment is difficult: the interested reader can refer to Prakken and Sartor (2008). Even more complex is to handle the interplay between the dialectical and the procedural layers (Prakken 2001).

16.3.1.4 Commitments

Norms about arguments are present also when we reason about commitments in dialogue (Brandom 1998). In particular, consider the case in which an individual is committed to a proposition. This means that she should only say things which she believes to be true. This is a norm about how to argue, and it poses constraints on the reasoning process. An individual is not permitted to first say that a proposition p holds, and later when she is questioned about the same assertion, simply to drop it. For instance, this kind of problems has been addressed by Boella et al. (2006). Future challenges in norms and argumentation applied to dialogues arise when we consider systems where more than one norm regulates the debate. How to detect the possible conflicts among these norms, and how to provide a mechanism such that conflicts are solved, are two open challenges for agreement technologies.

16.3.2 Epistemic Norms

Epistemic norms are norms that guide, regulate or control our epistemology (Fagin et al. 1995), that is, what should we know or believe, how should we acquire knowledge, and how should we know what we know? The well known AGM axioms of belief revision (Alchourrón et al. 1985) are thus epistemic norms, namely norms on how to change beliefs. For example, if you learn new information which you accept, then you should revise your knowledge in a minimal way to accommodate for this new information.

It is tempting to define epistemic norms as norms that involve an epistemic or doxastic operator, such as norms on what you should know, what you should believe, or what you are permitted to know. For example, in many legal systems there is a norm "you should know the law". Classically this was discussed by Åqvist (1967) in his paradox of knower, represented by the formula $OKp \rightarrow Op$ of traditional modal deontic logic: if you should know p, then p should be the case.

Castañeda (1988) developed the logic of epistemic obligation, Cuppens (1993) and Cuppens and Demolombe (1997) study obligations about knowledge in the context of computer security, and Aucher et al. (2010) further develop their logic and apply it to privacy regulations. Pacuit et al. (2006) study knowledge based obligations, such as the obligation of a doctor to help someone if he knows the patient is ill.

In the context of agreement technologies, and multiagent systems and computer science in general, epistemic norms lead to new challenges. We cannot look into the head of a human agent, and we therefore cannot verify whether a human knows or believes something. Consequently, we cannot verify whether a human agent complies with an epistemic norm that he is forbidden or permitted to know something. With artificial agents, we can often verify an agent program before it is accepted to a virtual organization, and we can verify that it conforms to the norms, or more generally complies with the epistemic policies. Likewise, we can let an artificial agent comply with epistemic norms by not making observations when the answer to the question might lead to a violation of an epistemic norm, or not ask another agent about some information in similar circumstances.

16.3.3 Enforcement Tools for the Burden of Proof

The notions of proof standards and burden of proof refer to argumentation as a dialogical process for making justified decisions. The process starts with an initial claim, and the aim of the whole process is to clarify the claim, and produce a justification of a decision. The process will return a set of claims together with the decision to accept or reject them. A fundamental part of the output process is the proof justifying the decision of each claim, to show how the decision is supported by the theory. As in legal reasoning, a proof in argumentation is a structure which

demonstrates to an audience that a claim satisfies its applicable proof standard. A summary of proof standards and burden of proof is presented by Gordon and Walton (2010).

The proof standards and burden of proof provide a kind of norm applied to argumentation. While several burdens of proof have been theoretically defined in the literature, such as *burden of claiming* and *burden of questioning*, a future challenge is the development of tools to support the humans operating in the legal field. The idea is to start from systems like Carneades,[3] which already provide a tool for modeling legal dialogues, and improve them to support the interaction with humans. For instance, a judge can use such a tool to look at the argumentation framework which models the trial, and she will be able to detect the possible "irregularities" with respect to the burden of proof. Moreover, the tool should provide the judge with a summary of the argumentation framework representing the trial's arguments. The same tool can be used by the lawyers to detect the possible weak points of a deliberation. In this way, the lawyer will know precisely which weak point to appeal. The development of such kinds of tools based on burden of proof poses a challenge in the area of norms and argumentation.

16.3.4 Conclusions

In this section, we presented an overview of existing work in the field of argumentation and norms. The analysis of this work allows us to identify various challenges of the domain. These challenges fall into two broad categories: how to argue about norms, and how norms influence the argumentation process. In particular, we highlighted the following challenges:

- Arguing about norms:
  1. *Societal modelling and control*:
     - Individuals debate about the merits of norms and its effects;
     - Individuals persuade others about utility of norm adoption;
  2. *Constitutive norms*:
     - More than two agents performing ontology alignment;
     - Avoiding the central ontology mapping repository;
  3. *Regulative norms*: considering norms in practical reasoning;
  4. *Normative constraints*:
     - Complex normative reasoning for deadlines, norm violation, norm fulfillment;
     - Using argument schemes to reason about norms being or not in force;

[3] https://github.com/carneades/carneades

5. *Normative conflict*: developing richer preference models and logics for reasoning about norm interaction;
6. *Practical reasoning*:
   - Integration of domain specific knowledge and inference using argument schemes;
   - New reasoning heuristics;
7. *Monitoring norms*:
   - Identifying argument schemes which reason about uncertainty;
   - Weighting up conflicting uncertain evidence;

- Norms about argumentation:
  1. *Dialogue*:
     - Interplay between dialectical and procedural norms;
     - Modelling dialogues where several norms regulate a dialogue;
  2. *Burden of proof*: tools for supporting people in legal field to verify proof standards;

These future challenges should be addressed both from the theoretical and the design point of view. The former involves defining new innovative models which integrate argumentation theory and norms, while the latter involves the creation of tools which leverage, apply, and implement various aspects of existing theoretical work.

References

Alchourrón, C.E., and D.C. Makinson. 1982. The logic of theory change: Contraction functions and their associated revision functions. *Theoria* 48: 14–37.

Alchourrón, C., P. Gärdenfors, and D. Makinson. 1985. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* 50(2): 510–530.

Åqvist, L. 1967. Good samaritans, contrary-to-duty imperatives, and epistemic obligations. *Nôus* 1: 361–379.

Artikis, A., M. Sergot, and J. Pitt. 2009. Specifying norm-governed computational societies. *ACM Transactions on Computational Logic* 10(1): 1–42.

Ashley, K.D. 1990. *Modeling legal argument – reasoning with cases and hypotheticals*, Artificial intelligence and legal reasoning. Cambridge: MIT.

Aucher, G., G. Boella, and L. van der Torre. 2010. Prescriptive and descriptive obligations in dynamic epistemic deontic logic. In *AI approaches to the complexity of legal systems (AICOL 2009)*, LNAI, vol. 6237, 150–161. Berlin: Springer.

Barwise, J., and L.S. Moss. 1996. *Vicious circles: On the mathematics of non-wellfounded phenomena*. Stanford: CSLI.

Bench-Capon, T., H. Prakken, and G. Sartor. 2010. *Argumentation in legal reasoning, argumentation in artificial intelligence*. Berlin: Springer.

Bex, F., H. Prakken, C. Reed, and D. Walton. 2003. Towards a formal account of reasoning about evidence: Argumentation schemes and generalisations. *Artificial Intelligence Law* 11(2–3): 125–165.

Billhardt, H., R. Centeno, C.E. Cuesta, A. Fernández, R. Hermoso, R. Ortiz, S. Ossowski, J.S. Pérez-Sotelo, and M. Vasirani. 2011. Organisational structures in next-generation distributed systems: Towards a technology of agreement. *Multiagent and Grid Systems* 7(2–3): 109–125.

Bochman, A. 2003. Collective argumentation and disjunctive logic programming. *Journal of Logic and Computation* 13(3): 405–428.

Bochman, A. 2005. Propositional argumentation and causal reasoning. In *Proceedings of the nineteenth international joint conference on artificial intelligence (IJCAI 2005)*, Edinburgh, 388–393.

Boella, G., and L. van der Torre. 2003. BDI and BOID argumentation. In *Proceedings of the third workshop on computational models of natural argument (CMNA 2003)*, Acapulco.

Boella, G., and L. van der Torre. 2004. The social delegation cycle. In *Proceedings of the 7th international workshop on deontic logic in computer science (DEON 2004)*, Madeira, 29–42.

Boella, G., and L. Van der Torre. 2004. Regulative and constitutive norms in normative multiagent systems. In *Proceedings of the ninth international conference on principles of knowledge representation and reasoning (KR 2004)*, Whistler, 255–266.

Boella, G., R. Damiano, J. Hulstijn, and L. van der Torre. 2006. ACL Semantics Between Social Commitments and Mental Attitudes. In *Agent communication II, international workshops on agent communication (AC 2005)*, Lecture notes in computer science, vol. 3859, 30–44. Berlin: Springer.

Boella, G., G. Pigozzi, and L. van der Torre. 2009. Normative framework for normative system change. In *Proceedings of the eighth international joint conference on autonomous agents and multiagent systems (AAMAS 2009)*, 169–176. New York: ACM.

Brandom, R. 1998. *Making it explicit: Reasoning, representing, and discursive commitment*. Cambridge: Harvard University Press.

Broersen, J., M. Dastani, J. Hulstijn, Z. Huang, and L. van der Torre. 2001. The BOID architecture: Conflicts between beliefs, obligations, intentions and desires. In *Proceedings of the fifth international conference on autonomous agents (AAMAS 2001)*, Ann Arbor, 9–16.

Castañeda, H.N. 1988. Knowledge and epistemic obligation. *Philosophical Perspectives* 2: 211–233.

Cuppens, F. 1993. A logical formalization of secrecy. In *IEEE computer security foundations workshop (CSFW 1993)*. Los Alamitos: IEEE Computer Society.

Cuppens, F., and R. Demolombe. 1997. A modal logical framework for security policies. In *Proceedings of the tenth international symposium on foundations of intelligent systems (ISMIS 1997)*, Charlotte, 579–589.

Daskalopulu, A., T. Dimitrakos, and T. Maibaum. 2002. Evidence-based electronic contract performance monitoring. *Group Decision and Negotiation* 11(6): 469–485.

dos Santos, C.T., P. Quaresma, R. Vieira, and A. Isaac. 2009. Comparing argumentation frameworks for composite ontology matching. In *Proceedings of the sixth international workshop on argumentation in multi-agent systems (ArgMAS 2009)*, Lecture notes in computer science, vol. 6057, 305–320. Berlin: Springer.

Dung, P.M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77(2): 321–357.

Fagin, R., J. Halpern, Y. Moses, and M. Vardi. 1995. *Reasoning about knowledge*. Cambridge: MIT.

Farley, A.M., and K. Freeman. 1995. Burden of proof in legal argumentation. In *Proceedings of the fifth international conference on artificial intelligence and law (ICAIL 1995)*, Edinburgh, 156–164.

Gabbay, D.M., G. Pigozzi, and J. Woods. 2003. Controlled revision – an algorithmic approach for belief revision. *Journal of Logic and Computation* 13(1): 3–22.

Gelati, J., G. Governatori, A. Rotolo, and G. Sartor. 2004. Normative autonomy and normative co-ordination: Declarative power, representation, and mandate. *Artificial Intelligence and Law* 12(1–2): 53–81.

Gordon, T.F. 1993. The pleadings game: Formalizing procedural justice. In *Proceedings of the fourth international conference on artificial intelligence and law (ICAIL 1993)*, Amsterdam, 10–19.

Gordon, T.F. 1995. *The pleadings game: An artificial intelligence model of procedural justice*. New York: Springer.

Gordon, T.F., and D. Walton. 2009. Legal reasoning with argumentation schemes. In *Proceedings of the twelfth international conference on artificial intelligence and law (ICAIL 2009)*, Barcelona, 137–146.

Gordon, T.F., and D. Walton. 2010. *Proof burdens and standards, argumentation in artificial intelligence*. Boston: Springer.

Governatori, G., and A. Rotolo. 2008. Changing Legal Systems: Abrogation and Annulment Part I: Revision of Defeasible Theories. In *Proceedings of the ninth international conference on deontic logic in computer science (DEON 2008)*, Lecture notes in computer science, vol. 5076, 3–18. Berlin: Springer.

Governatori, G., and A. Rotolo. 2010. Changing legal systems: Legal abrogations and annulments in defeasible logic. *The Logic Journal of IGPL* 18(1): 157–194.

Governatori, G., M. Maher, D. Billington, and G. Antoniou. 2004. Argumentation semantics for de-feasible logics. *Journal of Logic and Computation* 14(5): 675–702.

Haenni, R., B. Anrig, J. Kohlas, and N. Lehmann. 2001. A survey on probabilistic argumentation. In *Proceedings of the sixth European conference on symbolic and quantitative approaches to reasoning with uncertainty, workshop: Adventures in argumentation*, Boston, 19–25.

Horty, J. 1997. Nonmonotonic foundations for deontic logic. In *Defeasible deontic logic*, ed. D. Nute, 17–44. Dordrecht: Kluwer Academic.

Laera, L., I. Blacoe, V. Tamma, T.R. Payne, J. Euzenat, and T. Bench-Capon. 2007. Argumentation over Ontology Correspondences in MAS. In *Proceedings of the sixth international joint conference on autonomous agents and multiagent systems (AAMAS 2007)*, Honolulu, 228.

Makinson, D.C. 1998. On a fundamental problem of deontic logic. In *Norms, logics and information systems*, New studies in deontic logic and computer science, ed. H. Prakken and P. McNamara, 29–54. Amsterdam: IOS.

Makinson, D., and L. van der Torre. 2003. Permission from an input/output perspective. *Journal of Philosophical Logic* 32: 391–416.

McBurney, P., and S. Parsons. 2000. Risk agoras: Dialectical argumentation for scientific reasoning. In *Proceedings of the sixteenth conference on uncertainty in artificial intelligence (UAI 2000)*, Stanford, 371–379.

Medellin-Gasque, R., K. Atkinson, P. McBurney, and T. Bench-Capon. 2011. Arguments over co-operative plans. In *Proceedings of the first international workshop on the theory and applications of formal argumentation (TAFA 2011)*, Barcelona.

Modgil, S., and M. Luck. 2009. Argumentation in multi-agent systems. chap. *Argumentation based resolution of conflicts between desires and normative goals*, 19–36. Berlin: Springer.

Oren, N., T.J. Norman, and A. Preece. 2007. Subjective logic and arguing with evidence. *Artificial Intelligence Journal* 171(10–15): 838–854.

Oren, N., M. Luck, and T.J. Norman. 2008. Argumentation for normative reasoning. In *Proceedings of the symposium on behaviour regulation in multi-agent systems (BRMAS 2008) at AISB 2008*, Aberdeen, 55–60.

Oren, N., M. Luck, S. Miles, and T.J. Norman. 2008. An Argumentation Inspired Heuristic for Resolving Normative Conflict. In *Proceedings of the fifth workshop on coordination, organizations, institutions, and norms in agent systems (COIN@AAMAS-08)*, Toronto, 41–56.

Oren, N., M. Luck, and S. Miles. 2010. A model of normative power. In *Proceedings of the ninth international conference on autonomous agents and multiagent systems (AAMAS 2010)*, Toronto, 815–822. IFAAMAS.

Pacuit, E., R. Parikh, and E. Cogan. 2006. The logic of knowledge based obligation. *Synthese* 149(2): 311–341.

Pollock, J. 1995. *Cognitive carpentry*. Cambridge: MIT.

Prakken, H. 1993. A logical framework for modelling legal argument. In *Proceedings of the fourth international conference on artificial intelligence and law (ICAIL 1993)*, Amsterdam, 1–9.

Prakken, H. 2001. Relating protocols for dynamic dispute with logics for defeasible argumentation. *Syntese* 127: 187–219.

Prakken, H. 2005. AI & Law, logic and argument schemes. *Argumentation* 19(Special Issue on the Toulmin Model Today): 303–320.

Prakken, H., and G. Sartor. 1996. A dialectical model of assessing conflicting arguments in legal reasoning. *Artificial Intelligence Law* 4(3–4): 331–368.

Prakken, H., and G. Sartor. 2002. The role of logic in computational models of legal argument: A critical survey. In *Computational logic: Logic programming and beyond*, 342–381. Berlin: Springer.

Prakken, H., and G. Sartor. 2004. The three faces of defeasibility in the law. *Ratio Juris* 17: 118–139.

Prakken, H., and G. Sartor. 2008. A logical analysis of burdens of proof. In *Legal evidence and proof: Statistics, stories, logic*, ed. H. Kaptein. Aldershot: Ashgate.

Prakken, H., and G. Vreeswijk. 2002. Logics for defeasible argumentation. In *Handbook of philosophical logic*, vol. 4, 2nd ed, ed. D. Gabbay and F. Guenthner, 219–318. Dordrecht: Kluwer Academic.

Prakken, H., C. Reed, and D. Walton. 2005. Dialogues about the burden of proof. In *Proceedings of the tenth international conference on artificial intelligence and law (ICAIL 2005)*, Bologna, 115–124.

Rissland, E.L., D.B. Skalak, and M.T. Friedman. 1993. BankXX: A Program to Generate Argument Through Case-Base Research. In *Proceedings of the fourth international conference on artificial intelligence and law (ICAIL 1993)*, 117–124.

Searle, J.R. 1997. *The construction of social reality*. New York: Free Press.

Sergot, M.J., F. Sadri, R.A. Kowalski, F. Kriwaczek, P. Hammond, and H.T. Cory. 1986. The british nationality act as a logic program. *Communications of the ACM* 29(5): 370–386.

Stolpe, A. 2010. Norm-system revision: Theory and application. *Artificial Intelligence Law* 18(3): 247–283.

Toniolo, A., T.J. Norman, and K. Sycara. 2011. Argumentation schemes for policy-driven planning. In *Proceedings of the first international workshop on the theory and applications of formal argumentation (TAFA 2011)*, Barcelona.

van der Torre, L., and Y.H. Tan. 1999. Diagnosis and decision making in normative reasoning. *Artificial Intelligence and Law* 7: 51–67.

Walton, D., and E. Krabbe. 1995. *Commitment in dialogue: Basic concepts of interpersonal reasoning, SUNY series in logic and language*. Albany: State University of New York Press.