Introduction

Paul McNamara
Department of Philosophy, University of New Hampshire, Durham, USA

Henry Prakken
Department of Computer Science, Vrije Universiteit, Amsterdam, The Netherlands

1 Deontic Logic: problems and developments

We start with a brief overview of some main lines of research in deontic logic. Other recent overviews can be found in Meyer & Wieringa (1993) and Nute (1997a), while a thorough survey of fundamentals is provided by Åqvist (1984).

1.1 The emergence of Standard Deontic Logic

Deontic logic started as a special case of an already existing branch of philosophical logic, modal logic. The field was essentially inaugurated as an academic discipline by Georg Henrik von Wright, with his classical 1951 paper. Von Wright claimed that the deontic concepts of obligation, permission, prohibition, and indifference stand to one another in the same interdefinability relations as do the modal concepts of necessity, possibility, impossibility and contingency. For example, just as 'possible' can be defined as 'not necessary that not', he claimed that 'permitted' can be defined as 'not obligatory that not'. Furthermore, he claimed that they stood in other analogous logical relations to one another. For example, just as necessity implies possibility, obligatoriness was taken to imply permissibility. Of course, one other natural principle of modal logic has no sound analogue for obligation: the principle that necessity implies truth. Replacing the latter principle with the former in the normal modal logic T, resulted in the first mature modal deontic logic: the normal modal logic KD[1]. Widespread subsequent focus of this system by deontic logicians resulted in its being called 'Standard Deontic Logic (SDL)'. In SDL, the usual box and diamond of modal logic are replaced respectively by 'O' (read as 'it is obligatory that' or 'it ought to be that'), and 'P' (read as 'it is permissible that'). Assuming an otherwise standard modal syntax for the underlying language, SDL contains all propositional tautologies and the modus ponens inference rule, as well as these axioms

K. $\mathrm{O}(\varphi \rightarrow \psi) \rightarrow (\mathrm{O}\varphi \rightarrow \mathrm{O}\psi)$
D. $\mathrm{O}\varphi \rightarrow \mathrm{P}\varphi$
DfP. $\mathrm{P}\varphi \leftrightarrow \neg\mathrm{O}\neg\varphi$

and the rule of inference

RN. $\dfrac{\varphi}{\mathrm{O}\varphi}$

[1]Our system names are taken from Chellas (1980).

Sometimes, the definitions of $\mathrm{F}\varphi$ (φ is forbidden) as $\mathrm{O}\neg\varphi$ and $\mathrm{I}\varphi$ (φ is indifferent) as $\neg\mathrm{O}\varphi \wedge \neg\mathrm{O}\neg\varphi$ are also added.[2]

A derived inference rule of SDL is

$$\text{RM.} \qquad \frac{\varphi \to \psi}{\mathrm{O}\varphi \to \mathrm{O}\psi}$$

In words: all logical consequences of what is obligatory are also obligatory.

Von Wright himself proposed a slightly weaker system in two respects. Syntactically, because he saw the operators as applying to actions rather than propositions, he did not allow for embedded deontic operators in his language. And aximatically, he rejected the inference rule RN, since it validates $\mathrm{O}\top$, the obligatoriness of tautologies.

With the emergence of Kripke-style semantics for modal logics in the late 50s and early 60s, SDL received a standard possible-worlds semantic underpinning, the main features of which were these. From each world, w, there is always a non-empty set of ideal alternatives — worlds where things go as they should according to the norms of w. $\mathrm{O}\varphi$ is true in a world w just in case φ is true in all its ideal alternatives. And $\mathrm{P}\varphi$ is true in a world just in case φ is true in some of its ideal alternatives. The valid formulas on this semantic approach coincide precisely with SDL.

Another formulation of SDL has been proposed by both Kanger (1971) and Anderson (1958). They reduce it to alethic modal logic by defining 'it ought to be that' via a necessity operator and a propositional violation constant, v, standing for 'a violation of the norms has occurred': $\mathrm{O}\varphi =_{df} \Box(\varphi \to v)$. Among other things, this approach highlights a fundamental characteristic of norms, viz. that they can be violated.

1.2 'Paradoxes' of deontic logic

Much of the subsequent philosophical research in deontic logic has been driven by several so-called 'paradoxes' of deontic logic: problems or puzzles centering especially around SDL. We briefly discuss two of them.

The Chisholm Paradox

Probably the most discussed paradox is the 'Chisholm paradox' (Chisholm, 1963). It consists of three conditional obligations and one factual observation.

1 It ought to be that Jones goes to the aid of his neighbours
2 It ought to be that if Jones goes to the aid of his neighbours, he tells them he is coming
3 If Jones does not go to the aid of his neighbours, then he ought not tell them he is coming
4 Jones does not go to the aid of his neighbours

(3) is often called a 'contrary-to-duty obligation' since it regulates the violation of another obligation, viz. (1). The four sentences seem consistent, and each seems logically independent from the others. However, in SDL, all plausible formalizations violate one of these two desiderata. The most literal formalization

[2] Reading 'O' as both 'it is obligatory that' and as 'it ought to be that' is not completely uncontroversial, as is identifying indifference with $\neg\mathrm{O}\varphi \wedge \neg\mathrm{O}\neg\varphi$ on either reading of 'O'. See McNamara (1996).

$1'$ Og
$2'$ O$(g \to t)$
$3'$ $\neg g \to$ O$\neg t$
$4'$ $\neg g$

is inconsistent. By K and propositional reasoning $1'$ and $2'$ imply Ot, while by modus ponens $3'$ and $4'$ imply O $\neg t$; but D and DfP imply $\neg($O$t \wedge$O$\neg t)$, whence the contradiction. Moreover, two other obvious formalizations (both (2) and (3) as material implications of the form $\varphi \to$ Oψ or as 'large scope' implications of the form O$(\varphi \to \psi)$) violate the desideratum of logical independence: $4'$ implies $g \to$ Ot and $1'$ implies O$(g \to t)$.

Essentially, the paradox highlights the problem of the relation between two ways of deriving a conditional obligation from an unconditional one, the 'deontic detachment' of Oψ from Oφ and O$(\varphi \to \psi)$, and the 'factual detachment' of Oψ from φ and $\varphi \to$ Oψ. In SDL both are valid, causing the inconsistency of $1'$–$4'$.

Several strategies for solving this problem have been applied. One approach diagnoses the redundancy results as paradoxes of the material implication and not of the deontic operators, and replaces the material implication with a weaker conditional operator (e.g. Decew, 1981). However, to maintain consistency, this approach has to give up one of the detachment principles. Some can live with this (and usually reject deontic detachment) but others cannot, and therefore other approaches have been tried.

One way is to distinguish two kinds of obligations: deontic detachment yields something that should ideally be the case, while factual detachment derives what should be the case in the actual (often non-ideal) circumstances (Jones & Pörn, 1985).

Another way to reconcile the two forms of detachment is to restrict the validity of the detachment principles, notably to make them relative to the passage of time: deontic detachment holds until it has been settled (rendered unalterable) that the primary obligation (that Jones will provide aid) has been or will be violated, after which factual detachment holds and we can conclude that Jones ought not tell the neighbours he is coming. A problem with this solution is that it does not work for Chisholm's original example. Jones' merely deciding to not aid his neighbours while there is still time to tell is insufficient to settle, in the temporal sense that there is no future possibility, that he will not go to their aid: merely deciding to see to φ does not imply that φ thereby becomes unalterable. Accordingly, later this notion of settlement was refined to include a non-temporal deliberative kind of settledness, reflecting what an agent has decided to do (Carmo & Jones, 1997). Others have made the principles relative not to time or decision but to the available knowledge about violations, which brings an element of defeasibility into the analysis. (See further below.)

Besides the desiderata of consistency and logical independence, two other conditions for formalizations of the Chisholm set have been proposed. Jones & Pörn (1985) argue that (2) and (3) should have the same logical form (which rules out the above literal formalization), while Prakken & Sergot (1996) argue that the following 'pragmatic oddity' must be avoided. ($1'$–$4'$) imply that Jones ought to go to the aid of his neighbours and also tell them he is not coming, which seems odd.

Dyadic deontic logic

A technical device that is often used in modifications of SDL is a dyadic version of the deontic operators, viz. O(φ/ψ), which reads as 'it ought to be that φ given ψ' and P(φ/ψ), usually defined as $\neg$O$(\neg\varphi/\psi)$. Essentially, dyadic deontic logics interpret the normative 'if' as a special, already normatively-laden kind of 'if'.

This approach was introduced in Von Wright (1956) to deal with the implication paradoxes of his 1951 system. From the outset, there has been a great deal of controversy over the soundness of basic dyadic principles, and unfortunately there is no widely-studied basic dyadic system analogous to SDL (cf. Tomberlin, 1981).

The complexity of this controversy is beyond the scope of this brief introduction, so we merely provide a few indications. First of all, this approach usually endorses the following definitions:

$\mathrm{P}(\varphi/\psi) \leftrightarrow \neg\mathrm{O}(\neg\varphi/\psi)$	Def dyadic P
$\mathrm{O}\varphi \leftrightarrow \mathrm{O}(\varphi/\top)$	Def monadic O
$\mathrm{P}\varphi \leftrightarrow \neg\mathrm{O}\neg\varphi$	Def monadic P

Secondly, principles such as:

$\mathrm{O}(\varphi \to \psi/\chi) \to (\mathrm{O}(\varphi/\chi) \to \mathrm{O}(\psi/\chi))$	Dyadic K
$\mathrm{O}(\varphi/\psi) \to \mathrm{P}(\varphi/\psi)$	Dyadic D

are often endorsed, as are rules such as:

$$\frac{\varphi}{\mathrm{O}(\varphi/\psi)} \qquad \text{Dyadic RN}$$

However, one very important principle that is most characteristically rejected is the following:

$$\mathrm{O}(\varphi/\psi) \to \mathrm{O}(\varphi/\psi \wedge \chi) \qquad \text{Dyadic augmentation}$$

The rejection of this principle is usually motivated with examples of the following kind. Even if I should let you use my salt shaker, if you ask, it hardly follows that I should let you use it if you ask *and* I know it contains fatally poisonous material.

How about the detachment principles for dyadic O? Dyadic K above encodes the dyadic analogue of deontic detachment, whereas the dyadic analogue of factual detachment is:

$$(\mathrm{O}(\varphi/\psi) \wedge \psi) \to \mathrm{O}(\varphi|\top) \qquad \text{Dyadic factual detachment}$$

Note that if both of these detachment principles are endorsed, and in the Chisholm example (2) and (3) are formalized as $\mathrm{O}(t|g)$ and $\mathrm{O}(\neg t|\neg g)$, then the formalization is still inconsistent.

As for the semantics of dyadic deontic logic, an elegant approach is the one of Hansson (1969) and Lewis (1974)[3]. It assigns to a world w a preference ordering on all alternative versions of w. A formula $\mathrm{O}(\varphi/\psi)$ is true in w just in case φ is true in all most preferred ψ-alternatives of w. Dyadic augmentation is invalidated by this semantics since the most preferred ψ-worlds need not be the most preferred $\psi \wedge \chi$-worlds. Similarly, dyadic factual detachment fails on this semantics, but all other dyadic principles cited above are validated.

[3] As observed by Makinson (1993), this approach was also a forerunner of possible-worlds semantic of defeasible conditionals.

The gentle murderer

Another much-discussed deontic paradox is that of the gentle murderer (Forrester, 1984).

1 Smith ought not to kill his mother
2 If Smith kills his mother, he ought to kill her gently
3 Smith kills his mother

Intuitively, these sentences seem consistent, but in SDL the most natural symbolization of (2) and (3) yields 'Smith ought to kill his mother gently' by modus ponens but, since gentle killing logically implies killing, we obtain 'Smith ought to kill his mother' by RM, which contradicts 1 by D.

One difficulty of this paradox lies in the fact that, even though there is good reason to think that the Chisholm paradox and the gentle murderer reflect a uniform phenomenon centering on conditional norms, several solutions proposed for the Chisholm paradox do not help for the gentle murderer. For instance, temporal solutions are inadequate since both actions take place at the same time, and the distinction between ideal-actual does not seem to apply since both obligations seem actual (it seems hardly merely ideal that one should not kill one's mother).

One diagnosis of the gently murderer, inspired by Castañeda's (1981) approach to the paradoxes, is based on linguistic considerations. It consists in arguing that in (2) only the gentleness of the action is within the scope of the obligation operator, while the property of being murder falls outside its scope, so that (2) would not imply the (contrary-to-duty) obligation to kill. Other attempts to solve the paradox involve the rejection of ROK (e.g. Jackson, 1985), and yet other recent accounts are based on refinements of the dyadic approach to deontic logic (see Nute, 1997a for some examples).

With this fundamental background sketched, we next quickly highlight a number of other outstanding problems of deontic logic.

1.3 Other problems

Conflicting obligations and moral dilemmas

One philosophical problem of deontic logic is that of conflicting obligations. In SDL, they are ruled out, since it validates $\neg(\mathrm{O}\varphi \wedge \mathrm{O}\neg\varphi)$. However, many philosophers and deontic logicians have felt that this does not do justice to the nature of genuine moral dilemmas, which seem to involve two genuinely conflicting obligations (see e.g. several papers in Gowans, 1987). Curiously enough, the discussions of the Chisholm and gentle murderer paradoxes on the one hand, and moral dilemmas on the other, have remained completely separate: to our knowledge, nobody has ever proposed the consistency of conflicting obligations as a solution to these paradoxes.

Norms and truth

A recurring theme in philosophical discussions on deontic logic is whether a semantics of deontic logic can be based on a notion of truth. A positive answer violates the opinion held by many that norms are not 'truth-accessible' (subject to truth and falsehood). On the other hand, a negative answer must explain how the deontic operators can nonetheless be combined with the truth-functional logical operators (or why they cannot). An important distinction here is that of norms vs. propositions about the existence of norms, discussed in detail by Alchourrón (1969). According to him, only a logic of propositions about norms can be based on truth. Accordingly, Alchourrón & Bulygin (1981) have developed a logic of norms in which the deontic concepts are interpreted in terms of speech acts, in particular, acts of 'norming', such as commanding and permitting. The distinction has been employed for assessing the intuitive validity of certain deontic reasoning patterns. For instance, Alchourrón and others have argued that conflicting obligations are inconsistent in a logic of norms, but not in a logic of propositions about norms, since it should be possible to consistently describe an inconsistent system of norms.

Agency and time

Much research in deontic logic has been devoted to the interplay of deontic notions with such notions as time (e.g. Thomasson, 1981) and agency (e.g. Horty & Belnap, 1995). Some of this research was driven by the deontic paradoxes, but some of it has more general interest. For example, $\mathrm{O}\varphi$, can be read naturally as 'it is obligatory that φ' or as 'it ought to be that φ', but what of person-relative obligations such as my obligation to go to a meeting on tuesday, which are a central part of normative discourse? In order to do justice to personal obligations, we would seem to need to represent not only the person, but the agency of the person as well (for instance, her power to bring about some, but not other, states of affairs). Furthermore, if one is obligated to bring about a certain state of affairs, the state of affairs must lie in the future, and it would seem that it cannot be already settled now that it will occur. Since personal obligations are obligations that an agent has to act in a certain way or to bring about a certain state of affairs, one further philosophical problem with respect to agency is whether deontic operators apply fundamentally to actions or to propositions.

Defeasibility

In recent years it has been argued that several problems of deontic logic can be dealt with by embedding an account of deontic modalities in logics for nonmonotonic, or defeasible reasoning (see Nute, 1997a for a collection of papers on this topic). Nonmonotonic logics have been developed over the last 20 years in the field of Artificial Intelligence to model commonsense reasoning. Defeasible deontic systems have been especially useful in modeling notions made famous in ethics by Ross (1930) (prima facie obligations, one such obligation conflicting with, defeating or overridding another, actual obligations as non-overridden prima facia obligations, etc.) Such systems have also been applied in the analysis of moral dilemmas; see e.g. Horty (1994), Morreau (1996) and Prakken (1996). And recently, Nute (1997b), following Loewer & Belzer (1983), applied defeasiblity to Chisholm's paradox, arguing that deontic detachment is only defeasibly valid, in that it holds unless there is evidence that the primary obligation has been violated.

1.4 Philosophical applications of deontic logic

Some have used combinations of deontic and action logics to characterize the various 'normative positions' a person can have towards a certain action or state of affairs (Kanger, 1971; Lindahl, 1977), or to define the legal relations that can hold between people (Allen & Saxon, 1986). And deontic logic has been applied by Alchourrón & Bulygin (1971) to characterize various notions concerning normative systems, such as gaps, completeness, consistency, and the solution of a case. Furthermore, enriched deontic logics have been used to define other concepts of common sense morality, such as supererogatory acts, and to distinguish arguably distinct moral concepts, such as moral indifference from moral optionality, and what one morally ought to do from what one morally must do (McNamara, 1996).

2 Deontic logic, computer science and artificial intelligence

The connection between deontic logic on the one hand, and computer science and artificial intelligence (henceforth CS/AI) on the other hand is twofold. In one direction, computer scientists and AI researchers have discovered that various formal tools developed in CS/AI (e.g. nonmonotonic and dynamic logics) have interesting applications to traditional issues in deontic logic. For instance, as mentioned above, nonmonotonic logics have shed new light on the issues of prima facie obligations and moral dilemmas. Furthermore, dynamic logic, originally developed to formalize the behaviour of computer programs, has been used to define action-based deontic logics (Meyer, 1988).

In the other direction, AI researchers have used deontic logic as a tool in modelling legal reasoning. And computer scientists have discovered that computer systems (including their interaction with other computer systems and with human agents) can often be productively modelled as norm-governed. Jones & Sergot (1993) discuss these points in a general way. They claim that deontic logic must be taken seriously whenever in CS/AI it is necessary to make explicit, and then reason about, the distinction between what ought to be the case and what is the case.

We now give a very brief overview of uses of deontic logic in CS/AI. A more detailed overview can be found in Meyer & Wieringa (1993).

2.1 Applications to AI & law

A main branch of AI is knowledge representation, the effort to represent knowledge in a form that can be processed by the computer when solving a problem. Logic has been widely used for these purposes but, perhaps surprisingly, in the field of AI & law there are very few applications of deontic logic to the representation of legal knowledge. And those applications that do exist (notably McCarty, e.g. 1989), are still largely theoretical. One possible reason for this could be that many legal concepts are not deontic but either ‘quasi-deontic’ (‘negligent’, ‘responsible’, ‘tort’, ‘criminal offence’) or purely definitional (‘vehicle’, ‘nationality’). But another reason could be that deontic logic (and modal logic in general) does not have a long tradition in the study of computational issues, for which reason researchers in AI & law often resort to first-order computational tools. (But see the work of Artosi & Governatori and collegues, e.g. Artosi et al., 1996.)

2.2 Computer systems as norm-governed entities

Computer systems are dynamic entities, and we would like them to behave in certain ways. Ideally, one would design them in a such a way that the desirable behaviour is enforced; however, the actual world is not always ideal. Implementations of computer systems can deviate from their (ideal) specifications, or external factors may cause erroneous behaviour. Several computer scientist have found deontic logic a useful tool for specifying ideal behaviour of computers, for reasoning about the discrepancy between ideal and actual behaviour of norm-governed computers, and for directing system behavior when less than ideal conditions are detected. Khosla & Maibaum (1987) have defined an extension of modal action logic with deontic logic, for purposes of computer system specification. And Fiadeiro & Maibaum (1991) try to reduce deontic logic to temporal logic, for the same purposes. Dubois (1993) has also studied deontic logic for purposes of system specification.

Khosla & Maibaum also discuss the application of deontic logic to the specification of fault tolerant systems, i.e., systems in which undesired behaviour cannot be completely avoided. Deontic logic can be used here to specify the corrective actions that must be undertaken when a computer system somehow fails to perform as it should. Coenen (1993) also applies deontic logic to this issue.

Deontic logic has also been applied to the specification of soft integrity constraints for databases, i.e., constraints that should be satisfied by the database but that can as a matter of fact be violated. An example is a constraint on a library database that books ought to be returned within three weeks. (Wieringa et al., 1991; Kwast, 1993; Carmo & Jones, 1994).

2.3 Norm-governed interaction with computer systems

Interaction of humans with computer systems is often regulated. It is often desirable to exclude certain types of users from access to systems or files, or from performing certain actions with files (write, delete, execute). In many applications it is best to physically enforce these regulations, but other applications might require more flexibility, especially in contexts where users are reluctant to be too rigidly forced into a certain pattern of behaviour. In such contexts the main concern is detecting rather than preventing undesired behaviour; accordingly, the interaction policies can perhaps better be regarded as genuine norms, which can be violated. Deontic logic is a tool for specifying such norms. The specification of access control to computer systems has been studied by Cuppens, e.g. with Demolombe (1996), and by Jones & Sergot (1992), who characterize access policies in terms of the Kanger-Lindahl theory of normative positions.

2.4 Multi-agent systems

Another area of computer science where normative notions have come up is the new field of ‘multi-agent systems’. Here computer systems are regarded as relatively independent entities (‘agents’) that interact with other systems not according to some centralized design, but according to their individual characteristics. Thus it cannot be ensured that every agent acts in a desirable way, in particular not when agents have different goals or intentions. In analyzing this problem, notions have come up like negotiation, commitments, and social conventions. Normative notions are relevant here and deontic logic may be useful in analyzing them. See e.g. Dignum (1996) for general considerations on this topic, and Conte et al. (1998), who have used deontic logic in formalizing the conditions under which a rational agent will form the goal of following a norm.

2.5 Other possible applications

We end with suggesting two other areas of computer science where deontic logic might be relevant. One is that of *electronic commerce*, where commercial transactions are carried out by or with the help of computer systems (this area overlaps with that of multi-agent systems). These transactions are governed by commercial law, and deontic logic may be used in specifying the relevant provisions.

Another branch of computer science where normative notions play a role is *computer-supported cooperative work*. For instance, when the people involved have taken up commitments towards each other, a computer system might keep track of the various commitments and their fulfilment (as described by Winograd & Flores, 1986). And when the cooperation is governed by certain procedures, a computer system might check whether any procedural norms have been violated by the participants in an electronic meeting. The latter is another context where it might not be a good idea to make norm violation physically impossible: often a meeting (or other form of cooperation) is better served by small procedural violations, as we can all experience in our professional lives.

3 The papers in this anthology

The book starts with a contribution by the 'founding father' of deontic logic, *Georg Henrik von Wright*. He gives an 'autobographic' overview of the development of his thinking on several issues, such as dyadic formulations of deontic logic, norms and action, conditional obligations, permission, and the problem of deontic logic and truth. On the latter topic, von Wright sketches his various efforts to reconcile his conviction that norms are not truth-accessible with his opinion that it is still possible to define notions of normative consistency and entailment.

Norms and truth

Three other papers address the problem of deontic logic and truth. In his invited paper, *David Makinson* also denies the capacity of norms to be true or false, taking this denial to imply three things: (1) strictly speaking, there are no logical relations between norms but only between propositions about norms; (2) deontic logic must be a logic, not of norms, but of the status of norms in normative systems; (3) the construction of such a logic cannot appeal to some already given logic of norms. He attempts such a construction, building on a proposal of Alchourrón. The approach is neither semantic nor axiomatic in style, but rather is based on the idea of the iterative application of detachment to conditional norms, with consistency checks to avoid "drowning" of conditions that are contrary to duty.

Jaap Hage presents a philosophical basis for a similar view. Elaborating on ideas from analytic philosophy, he embeds his opinion that norms (defined as rules with deontic consequents) are not truth-accessible in a general naturalistic account of institutional and social rules. Rules cannot be true or false but can only exist or not exist, i.e., rules are not propositions but objects. Rules can be applied, which results in institutional facts, i.e., facts that owe their existence to the application of existing norms (for instance, a contract owes its existence to contract law). Normative propositions describe these deontic facts, and thus their truth conditions depend on the purely factual observations concerning the existence of a norm. The paper ends with a sketch of a possible semantics for this naturalistic account of normativity, one that bases the truth of a normative proposition at a world merely on the non-institutional facts holding at that same world.

Van der Torre & Tan also agree that norms are not truth-accessible, but they still think that a logic of norms is possible. They define such a logic in terms of 'update semantics', a recent approach to logical semantics, based on the idea that one knows the meaning of a proposition if one knows the change it brings about in one's information state. For norms, van der Torre & Tan turn this into 'one knows the meaning of a norm if one knows the change it brings about in one's preference ordering on worlds'. According to them, such a dynamic semantics captures the speech act character of a norm as an act of norming some behavior.

Agency and time

A number of the authors explore connections between deontic logic, agency and/or time. *Paul Bartha* uses a framework for integrating agency and a deontic operator, developed by Thomasson, Belnap, Horty and others. The agency component involves an indeterministic forward-branching temporal logic that allows one to represent the future possibilities open to an agent, as well as the things that the agent can bring about (her available choices), and the things the agent does bring about, in a given history. With a preference ordering placed on the histories, 'O' can be modeled in this framework, and thus what ought to be can be blended with agency. Within this framework, Bartha makes a number of distinctions (for example, between primary, secondary, conditional and contrary-to-duty obligations), proposes a solution to the Chisholm paradox and a strengthened version of it, and provides an analysis of defeasible obligations.

Mark Brown uses a similar framework. He takes future branches rather than moments as basic, to represent the actions and abilities of the agents, as well as related concepts of refraining and allowing. Then he adapts notions he has introduced elsewhere to this framework, including a distinction between two concepts of an obligation operator, that of an obligation whose fulfillment necessitates φ, and that of an obligation for which φ is both necessary and sufficient. Employing an enriched set of semantic notions, he defines a correspondingly rich set of diverse operators, and uses these to represent the fluctuation of, and interactions between the actions, abilities and obligations of multiple agents over time.

Jörg Hanssen examines two well-known systems of deontic logic, Aqvist's dyadic deontic logic G, and van Eck's system of quantificational logic with time-indexed dyadic deontic operators. Hanssen establishes an important mapping between G and a propositional variant of Van Eck's system, thus showing that these independently developed approaches are much more closely related than they seem to be on the surface.

Analysis of normative concepts

Lindahl & Odelstad develop a formal framework for understanding intermediate concepts — concepts that link one conceptual domain to a related conceptual domain. The motivation proceeds via legal intermediate concepts like 'ownership', which have been fruitfully seen as linking a set of legally recognized factual grounds (e.g. the various grounds sufficient for asserting that x owns y in a legal system) to a set of normative consequences (e.g. the obligations/rights x has as an owner of y in the legal system). The authors adapt tools from lattice theory to provide a deeper understanding of intermediate concepts by formally representing them as nodes linking one conceptual sub-lattice to another.

Paul McNamara examines the extent to which his earlier framework, 'Doing Well Enough' (DWE) for supererogation and related notions of common-sense morality can be fruitfully recast in an extended Andersonian-Kangerian modal framework. He presents DWE, notes some limits on recasting DWE logics, and identifies a significant subset of the DWE logics that can be fruitfully recast. He then proves an analog of the Fundamental Theorem for Canonical Models for normal modal logics, and employs it to generate a collection of determination results for the recast logics.

Claudio Pizzi first summarizes his earlier work employing an Aqvist-style revision operator and a strict conditional to generate a logical framework for a weak conditional. Then he uses this weak conditional to make a number of interesting distinctions between causal relations, along with proposals about how these relations map onto various pre-theoretical causal distinctions (causal relevance, causal preemption, causal overdetermination, concomitant causation, etc.). Pizzi ends with some speculations on how his various causal notions may figure in legal conceptions of responsibility and agency.

Defeasibility and norm conflicts

Several papers take up the subject of defeasible deontic reasoning, and three of them build on formalisms from Artificial Intelligence developed earlier by the same authors. *Donald Nute* defines a defeasible deontic logic on the basis of his well-known general 'defeasible logic'. With an eye to computer applications, the language of this logic is kept simple. Its main feature is a mechanism for handling contradictory conclusions: consistency can be restored by preferring one conclusion over the other. This is partly done with user-provided preference criteria, but conflicts between deontic and factual detachment (as in the Chisholm paradox) are automatically resolved in favour of the latter. Thus Nute provides an alternative type of restriction on deontic detachment, not in terms of time, but in terms of what is known about norm violations.

Bell & Huang's nonmonotonic logic combines possible-worlds accounts of belief, time, agency and preference. They extend it with a deontic part, which allows for the derivation of both prima facie and actual obligations. A prima facie obligation is also actual if it is jointly realizable with all other obligations that are as least as important. The temporal aspect of the logic allows obligations to change over time.

Cholvy & Cuppens discuss the problem of merging conflicting regulations, based on their earlier work on merging information from different sources. Regulations and their mergings are expressed by means of indices attached to deontic operators. When norms from different regulations conflict, an ordering on the regulations decides which norms 'survive' the merging. The authors provide an axiomatization and possible-worlds semantics for their system, but remark that the logic is only well-behaved under certain syntactic restrictions.

Becher et al. address Carlos Alchourrón's latest theory-revision account of defeasible conditionals. Among other things, they show that his definition of prima facie obligations as unconditional obligations derived from his defeasible conditionals, does not coincide with David Ross' well-known explication of this concept.

Finally, as remarked above, *Paul Bartha* also analyses the notion of defeasible obligation within his system.

Computer systems applications

Krogh & Jones identify a new domain of application for deontic logic in computer science, viz. the verification of cryptographic protocols, an aspect of computer security. Such protocols concern the transmission of encrypted electronic messages, and the exchange of keys for decoding the messages. One goal is to ensure that such a key can be securely distributed. A way to verify whether a protocol meets its goals, is to formalize it in logic and check whether the desired properties follow from the formalization. Krogh & Jones argue that this verification method can benefit from deontic logic, since they identify a class of protocol flaws that involves normative notions.

Firozabadi, Tan & Lee give a logical analysis of various fraud-like notions in terms of standard deontic logic combined with logics of belief and agency. The application that the authors discuss informally is the detection of fraud in electronic commerce. In this context, they view trade contracts as norm-governed systems coupled with detection systems for identifying violations of contractual obligations, as well as misinformation about their fulfillment.

Marek Sergot generalizes and further develops the Kanger-Lindahl theory on normative positions. This theory involves a method for generating all possible normative positions between (one or more) agents and an act-type. Among other things, Sergot adds to it the means to specify these positions at varying levels of detail, by a process of progressive refinement. He has implemented the resulting theory as a computer program which, Sergot hopes, can be of aid in specifying and disambiguating regulations, such as, for example, access control to databases.

Combinations of logics

Finally, a number of authors present work centering on combining different logics. *Gabbay & Governatori* argue for the applicability of a multi-modal indexed-modality approach to the complexity of normative systems. They sketch Gabbay's techniques for generating efficient completeness proofs for combinations of modal logics, including cases where the modalities are ranked by their labels, and advocate their promise for representing normative matters. For example, they discuss how norms from various regulations (such as constitutional, national and local laws) can be combined and their hierarchical relations expressed. In sketching the general background, they provide a number of basic metatheorems, including a detailed application of this semantic approach to the completeness of a particular labelled-tableau proof system developed elsewhere by Alberto Artosi and Governatori.

Louis Goble explores combining deontic logic with relevant logic. He suggests that this may be beneficial for two reasons. Firstly, dropping the D axiom of SDL (to allow for moral dilemmas) is less problematic in a relevant logic setting than in a classical approach. And secondly, replacing the material conditional with a relevant conditional avoids the logical redundancy problems in the Chisholm paradox. However, Goble argues that although a standard Kripke-style accessibility semantics for the deontic operator blends well with relevant logic, the widely endorsed alternative preference-based semantics for the deontic operator does not blend so smoothly with it, even though either semantic approach can be used unproblematically for SDL itself.

Krabbendam & Meyer explore the application of their previously developed 'release logic' framework to deontic logic. In particular, they show how two well-known approaches to deontic logic can be combined in this framework, viz. the Kanger-Anderson reduction of 'ought to be' to alethic modal logic, and Meyer's similar reduction of 'ought to do' to dynamic logic. They then argue for this approach as follows. Release logics allow for a relativization of necessity operators to context, and necessity operators are pivotal in both deontic components of their combined system: so their system can achieve a greater flexibility in representing the sensitivity of normative discourse to context.

References

Alchourrón, C.E. (1969). Logic of norms and logic of normative propositions, *Logique ét Analyse* 12, 1969, 242–268.

Alchourrón, C.E. & Bulygin, E. (1971). *Normative Systems.* Wien-New York: Springer Verlag.

Alchourrón, C.E. & Bulygin, E. (1981). The expressive conception of norms. In Hilpinen (1981), 95–124.

Allen, L.E. & Saxon, C.S. (1986). Analysis of the logical structure of legal rules by a modernized and formalized version of Hohfeld's fundamental legal conceptions. In A. Martino & F. Socci Natali (eds.), Automated Analysis of Legal Texts. Amsterdam: North-Holland Publ., 385–450.

Anderson, A.R. (1958). A reduction of deontic logic to alethic modal logic. *Mind* 67: 100–103.

Artosi, A., Governatori, G. & Sartor, G. (1996). Towards a computational treatment of deontic defeasibility. In Brown & Carmo (1996), 27–46.

Åqvist, L. (1984). Deontic logic. In D. Gabbay & F. Guënthner (eds.), *Handbook of Philosophical Logic*, Vol II, Dordrecht: Reidel, 605–714.

Brown, M.A. & Carmo, J. (eds.) (1996): *Deontic Logic, Agency and Normative Systems.* Workshops in Computing, London: Springer.

Carmo, J. & Jones, A.J.I. (1994). Deontic database constraints and the characterisation of recovery. In A.J.I. Jones and M.J. Sergot (eds.): *Proceedings of DEON'94: Second International Workshop on Deontic Logic and Computer Science*, Oslo, January 1994. Complex 1/94, Oslo: Tano Publishers, 56–85.

Carmo, J. & Jones, A.J.I. (1997). A new approach to contrary-to-duty obligations. In Nute (1997a), 317–344.

Castañeda, H.-N. (1981). The paradoxes of deontic logic: the simplest solution to all of them in one fell swoop. In Hilpinen (1981), 37–85.

Chellas, B. (1980). *Modal logic: An introduction.* Cambridge University Press, 1980.

Chisholm, R. (1963). Contrary-to-duty imperatives and deontic logic. *Analysis* 24: 33–36.

Coenen, J. (1993). Top-down development of layered fault-tolerant systems and its problems; a deontic perspective. *Annals of Mathematics and Artificial Intelligence* 9: 133–150.

Conte, R., Castelfranchi, C. & Dignum, F. (1998). Autonomous norm-acceptance. To appear in *Proceedings of the Fifth International Workshop on Agent Theories, Architectures, and Languages (ATAL'98).* Springer Lecture Notes in Artificial Intelligence, Berlin: Springer Verlag.

Cuppens, F. & Demolombe, R. (1996). A deontic logic for reasoning about confidentiality. In Brown & Carmo (1996), 66–79.

Decew, J.W. (1981). Conditional obligation and counterfactuals. *Journal of Philosophical Logic* 10: 55–72.

Dignum, F. (1996). Autonomous agents and social norms. Presented at the *ICMAS'96 Workshop on Norms, Obligations and Conventions*, Kyoto, Japan, December 10, 1996.

Dubois, E. (1993). Use of deontic logic in the requirements engineering of composite systems. In Meyer & Wieringa (1993), 125–139.

Fiadeiro, J. & Maibaum, T.S.E. (1991). Temporal reasoning over deontic specifications. *Journal of Logic and Computation* 357–395.

Forrester, J.W. (1984). Gentle murder, or the adverbial Samaritan. *Journal of Philosophy* 81: 193–197.

Gowans, C.W. (ed.) (1987). *Moral Dilemmas.* Oxford: Oxford University Press.

Hansson, B. (1969). An analysis of some deontic logics. *Nôus*, 3: 373–398. Reprinted in Hilpinen (1971), 121–147.

Hilpinen, R. (ed.) (1971). *Deontic Logic: Introductory and Systematic Readings.* Dordrecht: Reidel.

Hilpinen, R. (ed.) (1981). *New Studies in Deontic Logic: Norms, Actions, and the Foundations of Ethics.* Dordrecht: Reidel.

Horty, J.F. (1994). Moral dilemmas and nonmonotonic logic. *Journal of Philosophical Logic* 23, 35–65.

Horty, J.H. & Belnap, N. (1995). The deliberative Stit: A study of action, omission, ability, and obligation. *Journal of Philosophical Logic* 24: 583–644.

Jackson, F. (1985). On the semantics and logic of obligation. *Mind* 94: 177–195.

Jones, A.J.I. & Pörn, I. (1985). Ideality, sub-ideality and deontic logic. *Synthese* 65: 275–290.

Jones, A.J.I. & Sergot, M.J. (1992). Formal specification of security requirements using the theory of normative positions. In Y. Deswarte, G. Eizenberg & J.-J. Quisquater (eds.), *Computer Security—ESORICS 92*. Lecture Notes in Computer Science, 648, Berlin: Springer Verlag, 103–121.

Jones, A.J.I. & Sergot, M.J. (1993). On the characterization of law and computer systems: the normative systems perspective. In Meyer & Wieringa (1993), 275–307.

Kanger, S. (1971) New foundations for ethical theory. In Hilpinen (1971), 36–58.

Khosla, S. & Maibaum, T.S.E. (1987). The prescription and description of state based systems. In B. Banieqbal, H. Barringer & A. Pnueli (eds.), *Temporal Logic in Specification*. Lecture Notes in Computer Science, 398, Berlin: Springer Verlag, 243–294.

Kwast, K.L. (1993). A deontic approach to database integrity. *Annals of Mathematics and Artificial Intelligence* 9: 205–238.

Lewis, D.K. (1974). Semantic analyses for dyadic deontic logic. In S. Stendlund (ed.): *Logical Theory and Semantic Analysis*. Dordrecht: Reidel, 1–14.

Lindahl, L. (1977). Position and Change. A Study in Law and Logic. Dordrecht: Reidel, Synthese Library.

Loewer, B. & Belzer, M. (1983). Dyadic deontic detachment. *Synthese* 54: 295–318.

Makinson, D. (1993). Five faces of minimality. *Studia Logica* 52: 339–379.

McCarty, L.T. (1989). A language for legal discourse I. Basic features. *Proceedings of the Second International Conference on Artificial Intelligence and Law*, 180–189. New York: ACM Press.

McNamara, P.M. (1996). Must I do what I ought? (Or will the least I can do do?). In Brown & Carmo (1996), 154–173.

Meyer, J.-J.Ch. (1988). A different approach to deontic logic: deontic logic viewed as a variant of dynamic logic. *Notre Dame Journal of Formal Logic* 29: 109–136.

Meyer, J.-J.Ch. & Wieringa, R.J. (eds.) (1993). *Deontic Logic in Computer Science: Normative System Specification*. John Wiley & Sons, Chichester, 1993

Morreau, M. (1996) Prima facie and seeming duties. *Studia Logica* 57(1): 47–71.

Nute, D.N. (ed.) (1997a). *Defeasible Deontic Logic*. Dordrecht etc.: Kluwer Synthese Library.

Nute, D.N. (1997b). Apparent obligation. In Nute (1997a), 287–315.

Prakken, H. (1996). Two approaches to the formalisation of defeasible deontic reasoning. *Studia Logica* 57(1): 73–90.

Prakken, H. & Sergot, M.J. (1996). Contrary-to-duty obligations. *Studia Logica* 57(1):91–115.

Thomason, R.H. (1981). Deontic logic as founded on tense logic. In Hilpinen (1981), 165–176.

J.E. Tomberlin (1981). Contrary-to-duty imperatives and conditional obligations. *Noûs* 15: 357–375.

Ross, W.D. (1930). *The Right and the Good*. Oxford University Press, Oxford.

Wright, G.H. von (1951). Deontic logic. *Mind* 60: 1–15.

Wright, G.H. von (1956). A note on deontic logic and derived obligation. *Mind* 65: 507–509.

Wieringa, R.J., Weigand, H., Meyer, J.-J.Ch., & Dignum, F. (1991). The inheritance of dynamic and deontic integrity constraints. *Annals of Mathematics and Artificial Intelligence* 3: 393–428.

Winograd, T. & Flores, F. (1986). *Understanding Computers and Cognition. A New Foundation for Design*. Reading, MA. etc.: Addison-Wesley Publishing Company.