The Consultancy Game

Alexander Boer
Department of Computer Science and Law, University of Amsterdam
Email: aboer@lri.jur.uva.nl

Abstract. The Legal Consultancy Game is the game of giving legal advice professionally. It aims to be a description in general terms of the predictable patterns in an advice transaction between a consultant and a client. It is also a description of what the client demands of good legal advice. There are good reasons to describe consultancy as a game-like activity. Firstly, it shows how legal consultancy relates to legal argumentation. Secondly, it is attractive for very practical reasons to describe legal consultancy as a game that can be played by a computer. It provides us with a prescriptive theory of what legal information retrieval (IR) systems are trying to achieve through various technologies.

1. Introduction

The Legal Consultancy Game is the game of giving legal advice professionally. It is, or aims to be, a description in general terms of the predictable patterns in an advice transaction between a consultant and a client. It is also a description of what the client demands of good legal advice. Like in any game there are rules that govern which states the game can be in, and which transitions between states are allowed. There is also a criterium to determine in which states the game ends. It is admittedly a bit of a strange game, because there is no real winner as such, but that never stopped Wittgenstein and many others after him from describing conversation in general as a game-like activity (cf. [9]). There are good reasons to describe consultancy as a game-like activity. Firstly, it shows how legal consultancy relates to legal argumentation – a veritable game-like activity with a real winner. Secondly, it is attractive for very practical reasons to describe legal consultancy as a game that can be played by a computer. It provides us with an prescriptive theory of what legal information retrieval (IR) systems are trying to achieve through various technologies. This paper attempts to bridge the gap between IR and argumentation.

The performance of IR technology is typically operationalised in terms of its ability to retrieve the – elements of – documents that are relevant to the information needs of the user community that uses the technology. The effectiveness of IR is often measured in terms of precision and recall [2,6]. Precision is the ratio of relevant document references retrieved to all references retrieved. Irrelevant document references that are retrieved (false positives) therefore lower precision. Recall is the ratio of relevant document references retrieved to all relevant references available. Failure to retrieve relevant document references (false negatives) therefore lowers this ratio. Given the potential importance of missing relevant document references, legal IR services cannot easily sacrifice recall in favor of precision.

Evaluation criteria – recall and precision – require an a priori distinction between relevant and irrelevant document references. In empirical investigations this distinction is often taken for granted and elicited directly from the prospective user community (or worse – researcher's private opinions). The danger of this practice is that it makes it impossible to

Alexander Boer, 'The Consultancy Game', in: Joost Breuker, Ronald Leenes and Radboud Winkels (eds.), *Legal Knowledge and Information Systems. Jurix 2000: The Thirteenth Annual Conference*. Amsterdam: IOS Press, 2000, pp. 99-111.

generalize precision and recall of some IR technology to other document samples and other user communities than those involved in the investigation. The most important tool for maximizing recall in legal IR is therefore a good criterium – one as immune to the idiosyncrasies of individuals and organizations as possible – of what document references are relevant to a legal problem. This paper proposes that a document reference is relevant if it has some definite role in the solution to the problem that made the user consult the IR service. This notion is only useful if it is possible to identify predictable patterns in the questions from an interesting user community. These patterns are to be found in an analogy of consulting legal IR services with legal consultancy. IR systems that index legal documents play a game of legal consultancy and a document element is relevant to a question if it used as a warrant in a valid argument in the answer to the question.

2. The Aboutness Problem in Legal IR

IR services for legal documents often depend on generic technology – technology that is not specific to a legal audience or legal documents (cf. [24]). The most common method currently used in commercial systems for legal IR is text search. Its great advantage is that little or no design and subject indexing effort is needed for specific domains. And some would even argue that the meaning of a document is best preserved anyway if we use the full text as it stands as its subject index. Certainly in a legal context this argument is problematic: The text itself is an imperfect description of an underlying thing – the Law. Even though we cannot agree on the nature of that underlying thing, we are inclined to believe that it is possible to translate, rewrite, or edit the document that describes the Law, without changing how that Law is applied in specific cases. We also believe that we may not find the document if we need it because we searched for the wrong paraphrases, even though we had the right concepts in mind.

To remedy this problem legal documents are often explicitly indexed and categorized by a human indexer. The selection of subjects and the categorization is however highly subjective and is often specifically targeted at the requirements of specific user communities – often it addresses legal professionals by specialization. It easily misses out on important aspects of the aboutness of documents. To overcome this aboutness problem we continually develop more advanced IR technologies and indexing methods to find implied documents [1,7,15], irrespective of the exact paraphrases used. How well an IR technology performs as a problem solver remains a function of how good its subject index is, and there is no easy, objective way to measure that if it is not sufficiently clear what the ideal answer to a query is. This is the key problem of *subject indexing*.

Technological advancement in IR technology is not likely to change the nature of the problem: The rapid proliferation and change of electronic legal documents and the increasing diversification of user communities with unique demands leads to an ever stronger demand for quality IR. Automated indexing approaches to IR (learning, statistics, neural networks) eventually depend on direct feedback on what document references are and are not relevant to the question. High quality IR will arguably always require interpretation of a document's subject matter by prospective users, because only those users will understand the contexts of usage in which (non-lexical) paraphrases can constitute a useful concept [22]. This indexing effort is – to put it bluntly – costly, and should be approached in a methodical way.

Surely, the number of strictly *lexical* concepts, paraphrases with a fixed meaning, currently in use is finite (but always changing) and a significant slice has been captured and categorized in lexical knowledge bases like WordNet [12]. But who knows what "independently driven" means? Learn one meaning later in this paper. The meaning of this

phrase derives from the existence of a specific user community who recognize a context and know what it means in that context. Off course, whatever indexing can be done automatically with linguistic knowledge (as in [1, 24]) is a bonus, and leaves us more time for more the exciting indexing problems.

The subject indexing problem is at the very core of library and information science. Coherent general theories on the purpose of indexing are however not available [13, 22] in either science. The questions are the same for all subject domains: Why is a given subject representation better than another? How does one create a good subject index in a systematic and controlled way? How does one determine the quality of a given subject index? The solutions are apparently different for each user community. Empirical investigations into the possibility of intersubjectively replicable document indices are often either too contrived to be really indicative of anything, or based on user communities or document samples that are not actually interesting [22]. They can be misleading as well: Apparent inconsistencies are actually often a natural outcome of different indexing contexts and purposes, each leading to alternative, valid indices for the same document. That subject indices are difficult to compare is a partly a consequence of the perceived indexing purpose: Representing knowledge that distinguishes elements in the domain and only representing it at the level of detail required for the IR technology.

It is quite fashionable in business process re-engineering (BPR) to advocate special-purpose IR interfaces and indices for user communities as distinguished by task in an organization (legal or otherwise). This practice makes user's information needs more explicit, but also limits the validity of indices to very specialized, and maybe temporary, user communities. But legal information scientists can do better than that: An alternative approach to legal IR is to explicitly aim to emulate the dialog between a general, naïve, unspecialized client solving the generic *legal* problem and a consultant. This generic problem context is legal consultancy. Describing it does not decrease the indexing effort; It provides us with an interpretation of what the question, formulated in an interface that is poor compared with natural language, means, and with a standard to compare the answer with.

3. What is Legal Consultancy about?

The layman's conceptual view on Law is quite straightforward; It divides legal knowledge into two orthogonal types of information: rules and cases. The pivotal legal problem in this view is to assess whether a case complies with (or deviates from) the rules. Rules (or norms) are typically interpreted as deontic formulas that allow, require or prohibit certain categories of cases, and cases as constellations of situations or events potentially regulated by rules. Written Law and knowledge representations that represent Law are intensional descriptions of what is allowed, required, and prohibited. Historical information about the application of Law is the extensional description of the same stuff. The intensional perspective is exemplified in AI & Law by i.a. deontic logics, and, more recently, generic and reusable task models and ontologies for legal assessment, e.g. in [3,5,26,29]. The latter approach, coined Legal Information Serving (cf. [3]), aims to provide KADS-like blueprints for the development of advanced legal IR systems that have a strong problem solving flavour. In some domains related problems like e.g. prediction of sentences, or determination of a causal relation between a person's actions and some (supposedly illegal) event, take the main stage. For most layman users of a legal IR system, however, the goal of seeking advice on the law is to be able to comply with it. To drive a car, for instance, it suffices to understand what is, and what is not, allowed in traffic. Some drivers may know what punishment awards which violation, but such knowledge is not at all required to drive.

Legal IR often focuses on historical information (Case Bases) instead of written law. The goal of a game of legal consultancy presented as in this paper is to help clients comply with the rules that regulate their lives. All client queries are interpreted in this context.

To describe or discuss this game we do not really need to know much about law, we only need to know how it is applied. Deconstructivist Schlag [21] cynically remarks that the games of law can be played (and won) even if one doesn't know what the ball (the law) one plays with looks like. Indeed, the one thing that can be counted upon is that no one engaged in "doing law" will ever raise the question. Schlag proposes that the following rhetorical hierarchy is what actually guides legal practice:

1. Do not confront an ontological question if it can be handled as an epistemic question.
2. Do not confront an epistemic question if it can be handled as a normative question.
3. Do not confront a normative question if it can be handled as a technical question.

This is an interesting perspective on the troublesome relation between legal theory and practice. By and large, neither judges nor any other bureaucratic decision makers listen to academic advice, because it always appears to address the wrong kinds of questions. Unfortunately, an information scientist must also address the ontological and epistemic questions first: The contents of the legal document must be represented, after all, and the indexing language has to be fixed before actual indexing starts. The ontological choices one makes are to a large extent pragmatic and open to criticism individually. The name of the game that is played by the computer, however, remains the same.

From a legal consultancy point of view, legal documents are about what ought and what ought not to be the case, and about who is held responsible for violation. The legal document is more than that, and has additional meanings to more specific user communities. But this is the baseline. Several theories have been developed in legal information science that address the overall requirements for automated assessment of simple legal queries. The FOLaw of Valente [26] and CLIME [27,29,30] frameworks distinguish three basic types of information structures:

1. The query is a cube, or conjunction, of facts; The case. Because the scope of the queries is limited to assessment, the exact phrasing or focus of the question asked is irrelevant.
2. Legal documents contain norms (obligations, permissions, or prohibitions) – functions that classify categories of cases as allowed or disallowed. These norms are ranked in order of priority to resolve conflicts between contradicting classifications by two norms of a single case. This priority ordering reflects certain concepts from legal theory that help solve apparent contradictions in written law.
3. Legal documents also require a monotonic (truth-preserving) theory about the regulated world used to infer additional facts implied by the case. This knowledge is intended to translate the phrasing used in the query to the phrasing in written law. Some of this theory may be explicit in the document.

Depending on the characteristics of the domain standardised case frames [28] or dynamically generated dialogs (cf. [17,15]) can be used to elicit the facts that make up the case. This theory uses monotonic rules to represent the world theory that relates categories of cases. The case to be assessed must be sufficiently complete to be correctly classified. Other theories, those aimed at representing argumentation (cf. [19]), do not distinguish norms and world theory (cf. [4]), and represent both with nonmonotonic rules to reflect contradictions in written law. This alternative representation is more expressive in this respect, but makes less commitment to what is modelled and is more difficult to use

effectively for subject representation by legal experts. Much of this work [18,19,20] as such appears to aim more at a argumentation-based logic programming language (or application programmer's interface) for legal purposes in this sense. On an ontological level, however, these theories are valid alternatives for representing legal consultancy scenarios. In fact, good argumentation is essential to good advice. This paper is only marginally concerned with the particulars of representation and argumentation – it suffices to be able to distinguish a valid argument from an invalid one. Instead it focuses on what document elements are required to help the user of an IR system solve the question by himself.

4. An Example of Argumentation

In argumentation frameworks like MacKenzie's DC, Gordon's Pleadings Game, and Lodder's DiaLaw (cf. [11]) there is a case to be argued and two players; A *plaintiff* and a *defendant*. Following DiaLaw, a proposition can be claimed, questioned, accepted, or withdrawn. The burden of proof for any proposition always rests with the player who claimed it. Argumentation Frameworks aim to model the dialog that ensues between plaintiff and defendant in the actual assessment of a case in terms of information states in the mind of plaintiff and defendant. and constraints on transitions between those states. This example considers a very restrictive argumentation game that fits with the case and rule representations in CLIME, translated into natural language for readability. Consider the following case description:

On a passenger ship, there are two fire pumps. The main engine drives one fire pump. The other fire pump is driven by an auxiliary power source.

Assume that the only relevant rules regulating this case are in table 1. Assume that plaintiff initially accepts the burden of proving that the case is disallowed, while the defendant accepts the burden of proving the opposite. Also assume that any case is implicitly allowed, unless the opposite is decided, and that therefore the plaintiff always takes the first turn. The following somewhat contrived dialog may develop in an argumentation framework:

Table 1: Example rules used throughout this paper

R1:	On a ship, there ought to be a fire pump.
R2:	(a) On a passenger ship, there ought to be at least two fire pumps. (b) One of these pumps may be driven by the main engine.
R3:	Fire pumps are to be independently driven.
R4:	If ballast or general service pumps are used as fire pumps, they are not to be used for pumping flammable liquids.
R5:	If fire pumps are subject to occasional duty for the transfer of fuel oil, change-over arrangements are to be fitted and so arranged as to prevent any mishandling.

1. Plaintiff – "The fire pump that is not independently driven is not allowed."
2. Defendant – "Why?"
3. Plaintiff – "A passenger ship is a ship. A fire pump driven by the main engine is not independently driven. Fire pumps on a ship are to be independently driven, according to R3."
4. Defendant – "I agree, but the fire pump driven by the main engine is nevertheless allowed."

5. Plaintiff – “Why?”
6. Defendant – “The ship is a passenger ship and there is a second fire pump that is independently driven, because it is driven by an auxiliary power source. According to R2(b) it is allowed that on a passenger ship one fire pump is driven by the main engine and another fire pump is independently driven. Rule R2(B) is an exception to R3.”
7. Plaintiff – “I agree. It is allowed.”

The conclusion is plausible and the same one assessment in FOLaw and CLIME would reach in different ways. The game starts when the plaintiff consults the rules in table 1 and constructs an argument to support the claim that the case is disallowed. In turn 2 the defendant questions this claim. In CLIME terminology: Of the three additional claims the plaintiff makes in response in turn 3, the first two claims introduce facts derived from the world theory, the third one introduces a norm and a reference to document element R3 to warrant it. That e.g “independently driven” in this context means that it is not driven by the main power source of the ship requires the theory about the regulated world. The original claim is now entailed by the case and warranted. In turn 4 the defendant consults the rules in table 1 and accepts all three claims made by the plaintiff but rebuts the original argument with a claim that directly denies the original claim. The plaintiff questions this claim in turn 5. The defendant supplies three additional claims in turn 6, introducing some world theory, a norm and a reference to rule R2(B) to warrant that norm, and a priority ordering over these norms including a reference to a principle (“an exception”) to warrant it. The norm that is introduced disaffirms the norm introduced by the plaintiff and is therefore prior in case of conflict. The concept of disaffirmation will be defined later. The plaintiff consults the rules in table 1 and fails to construct any counter arguments. He therefore accepts the three claims, withdraws the claim that the case not allowed and accepts the claim that it is allowed and the defendant wins.

In this case the course of dialog is almost inevitable, and it may seem that are no other equally efficient dialogs. This is not the case in general: Remove the pump driven by the auxiliary power source from the case and there are choices for the plaintiff. As information is added to the case or the rules in table 1, the number of choices for plaintiff and defendant, and the number of turns needed to reach a conclusion will typically increase. This example employs the algorithmic strategy used by CLIME (barring defeasible inferences from the theory about the regulated world) to minimise the number of turns and warrants required to come to a final classification. This strategy is not typical of argumentation as we find it in correspondence between legal professionals: It is an idealized strategy.

5. Argumentation for Automated Legal Consultants

To change the argumentation game to a question-answering game, a Consultancy Game, we introduce a third player, the Client, in need of legal consultancy. Plaintiff and defendant, now played by one schizophrenic legal expert – the consultant, have no direct access to the facts of the case and start with the information that can be obtained from the client’s question. The client, who has no way of guessing what information is relevant to the decision, may for instance pose the following question on the phone:

Client – “Is it allowed to use the main engine as a power source for a fire pump on a ship?”

Note that the focus of the question is completely irrelevant. The consultant simply reacts in accordance with his roles; the plaintiff tries to argue that the case implicit in the question is not allowed, while the defendant tries to argue it is allowed. If having a fire pump on a ship would be disallowed in itself, the plaintiff persona would co-operatively point that out. The strategy that seemed a bit contrived in the argumentation example in the previous section is now more realistic: Plaintiff and Defendant are now played by the same legal expert and therefore share the same knowledge base. The strategy is typical of for instance phone conversations between a legal advisor and a client about topics that are routine to the advisor. As the game starts the expert in his role as plaintiff consults the rules in table 1 and claims that the case is disallowed:

Plaintiff – "A passenger ship is a ship. A fire pump driven by the main engine is not independently driven. Fire pumps on a ship are to be independently driven, according to R3. A fire pump that is not independently driven is therefore not allowed."

The above is a single argument pleading disallowed in a single turn of the game. The defendant persona has no information to work with, but he learns from table 1 that there is a norm that would attack (rebut) the argument brought forward by the plaintiff. The conditions of applicability of this norm are consistent with the information in the question. If however no additional information can be obtained from the client, the defendant has no options left and the client may be left with a misleading answer to his question. This problem is a consequence of incomplete input. If the defendant is however able to ask yes/no questions of the client, input completeness is restored and the following dialog solves the problem:

1. Defendant – "Is this ship you consider a passenger ship?"
2. Client – "Yes."
3. Defendant – "Is there a second fire pump on this ship?"
4. Client – "Yes."
5. Defendant – "Is this second fire pump independently driven?"
6. Client – "Yes."
7. Defendant – "According to R2(B), it is allowed that on a passenger ship one fire pump is driven by the main engine and another fire pump is independently driven. Rule R2(B) disaffirms R3. The fire pump driven by the main engine is therefore allowed."

The above exchange is a single argument pleading allowed in an underlying argumentation dialog. The defendant wins the dialog because the plaintiff cannot reasonably complete a valid turn with the available rules without changing the subject of the dialog. There are no other new arguments that follow from the facts client and expert both know. Observe that the client can add claims about the facts of the case, first in the original question and later in response to questions from plaintiff and defendant. Assume in addition that:

1. The client can make factual claims but is restricted by a user interface from inserting contradictory facts – the assumption of input correctness, a requirement for a monotonous inference system.
2. Plaintiff and defendant can claim by defining or applying world knowledge, norms, or ordering principles.

3. The client cannot claim by defining or applying world knowledge, norms, or ordering principles because he lacks the knowledge and authority to do so. Plaintiff and defendant cannot question the facts presented by the client, because they do not have independent access to the facts.

An IR system capable of handling the example question of this section should return R2 and R3 to maximise recall, and it should not return R1, R4, and R5 to maximise precision. All of them are on the topic of 'fire pumps', but that does not distinguish them from the rest of the chapter they are in. If the system claims to rank the answer in order of priority, R2 should be prior to R3, because the user is entitled to stop reading after finding R3, since it appears to contain the answer. This ranking is static and does not require knowledge representation beyond the priority ordering itself. An IR system is able to represent the question if the user can express the constellation of facts implicit in the question: There is a ship that has two parts, a fire pump and a main engine, and the main engine is the power source of the ship.

The Consultancy Game is a conceptual framework for IR systems, subject indexing, and the design of associated indexing tools. Legal IR systems do not start with complete subject representations – these are after all created by experts capable of legal assessment. The Consultancy Game makes no distinction between dialogs with the client intended to solve a problem and dialogs intended to acquire arguments from experts in the role of plaintiff and defendant for subject representation. End users and experts can be directly interfaced (with proper restrictions on interfaces) as long as the domain of interest is not sufficiently covered. The legal IR system can assume the role of plaintiff and defendant when communicating to a client if able, and of the client (asking hypothetical questions for knowledge elicitation) when communicating to an expert. In the latter case, one uses a specific strategy designed to cover gaps in the subject representation, but it plays the same game.

6. Document Indexing Considerations for Consultancy

What does this the Consultancy Game tell us about how legal hypertexts in HTML or XML should be linked up? A legal document may regulate a case in a number of ways, but by far most important are norms that allow or disallow a category of case. A rule in the document that presents a case but is clearly not a norm, is often a scope restriction. Let's take an example that is classified by Prakken and Schrickx in [20] as an implicit exception; Section 2 of the Dutch Rent Act (HPW) states that the code is not applicable to rent contracts which by their nature concern a short term usage. This rule applies to the case in which there is a rent contract that does not concern short term usage (which can be expressed as a case). It also states that this case to which it applies subsumes all cases in norms that refer to the HPW. This relation over document elements refers to a wholly different domain – that of the structure of the document.

There are two sources of nonmonotonicity in application of rules to a case that can be abstracted to links between document elements. Firstly, a rule may have priority over another rule because of principles like *lex posterior* or *lex superior*. Priority is independent of the contents of the rule: It is determined by attributes of the legal document itself. The ordering is in reality on norms instead of document elements because a document element may have multiple positions in the priority ordering as it applies to different cases – but it is practical to prioritise the document elements directly and forget about the norms. If a norm is applied in the active turn of the game, no norm with less priority has to be considered in subsequent turns. Secondly, a norm expressed in the rule may be an exception to another

norm in another rule. This is commonly called the *Lex Specialis* principle. These exceptions provide the ammunition for attacks on claims in argumentation.

If the inference system is partly monotonic, like CLIME [27,29,30], the possibilities for attack are limited compared nonmonotonic systems like e.g. [19]. The arguments subject to attack are the arguments asserting the applicability of a norm. There are basically two ways to construct an argument that completes a valid turn.

Firstly, the case in the active turn may match a case regulated by a norm that is classified in accordance with the role (plaintiff or defendant) of the player who has to make a turn. The player accepts the currently active argument, if any, but introduces a new argument. This strategy is taken by the defendant in the first move. Consider the following variant question by the client:

Client – "Is it allowed to use the main engine as a power source for a fire pump on a ship, if the fire pump is subject to occasional duty for the transfer of fuel oil?"

The example dialog still fits the question, but now the plaintiff can make a move after the defendant rebutted his first argument. The plaintiff shifts the attention to the transfer of fuel oil and questions the changeover arrangements to verify the argument that this is not allowed according to R4.

Secondly, there may be a valid rebutting attack (cf. [16]) on the last argument, as in the example game, that *disaffirms* the norm applied in the active turn of the dialog through *Lex Specialis*. This rebutter can be used to ask questions of the user to counter the active argument. The turn is valid if the user confirms all questions. This strategy is preferred. Since there are many such rebutting cases for any norm and there are implications between them, it makes sense to pick the best, or minimal, rebuttal. A purely relevance based way of computing the best rebuttal, the one that is most likely to be first in the smallest sequence of questions required to rule out all possible rebuttals, is based on a generic method for picking the most relevant question in a decision procedure (cf. [23]). This calculation is closely related to the prime implicant required for the match between cases [26], except that in this view it can be performed offline because it only refers to two norms. If the knowledge representation is quite simple, this procedure need not be prohibitive. Recent algorithms for enumerating prime implicants of a CNF or DNF formula are discussed in [14].

If the system is nonmonotonic like e.g. [19], two other options are available, both qualified as undercutting attacks (cf. [16]). The third option is that an inferred fact supporting an argument may be undercut with a contradictory and prior inference. This is a consequence of nonmonotonous reasoning with the theory about the regulated domain. Fourth option is introduced by the possibility of applying a norm to the case, even though there is a prior norm that also applies. Pointing out that the other norm is prior and equally valid completes a turn.

The CLIME [27,29,30] and FOLaw [26] frameworks introduced earlier make some pragmatic choices on representation: The theory about the regulated domain is monotonic (truth-preserving) and the beliefs of plaintiff and defendant are closed under logical consequence. They must resolve any inconsistency that arises through insertion of a new definition immediately to maintain input correctness.

It is interesting to include as a corollary that the user cannot withdraw arbitrary facts from the case but he can always roll the dialog back to the previous turn. The set of facts to which the user is committed behaves in effect as a stack; The fact that was pushed on the stack last, is popped first.

The result is that if plaintiff and defendant have the competence to ask the right questions of the user and the beliefs of plaintiff and defendant concerning the facts are closed under logical consequence, then the plaintiff has a good winning strategy if the case considered by the user is disallowed, and the defendant has a good winning strategy if the case is allowed. The thing to do is to use any available new argument and rebutting attack available to present a complete answer and maximally relevant set of document elements. This strategy was used in the contrived example earlier in this paper. Another advantage of monotonicity for IR technology deployed on the internet is that the server has a great degree of freedom in how much information is stored between turns (minimally a fact stack and a warrant – or applied rule – stack where the warrant of the active argument is on top), so that available memory and inference resources can be used optimally.

For practical IR indexing it is important to realise that the hard parts, computationally speaking, are computation of the first match of the case to a norm to construct new arguments, and subsequent calculation of the right questions to ask, iff there is any ambiguity – which is rarely the case. The first match of the query to a document element can be approximated with any IR technology. There is no real need to calculate the right questions. Indexing concepts and discovering possible attack links between document elements is enough to meet the standard set by the Consultancy Game. Structural attack links between norms are quite simple and all static. If one document element matches the question, other documents are immediately implied. The four most interesting to be discovered during indexing are affirmation, disaffirmation, disaffirmation conflict, and compliance conflict.

Affirmation – A norm affirms another norm if it classifies a more specific case, with an equal classification (both allowed or both disallowed). A norm that affirms another norm is logically redundant, but quite common in legislation for technical domains. A classification society or a national government, for instance, may design rules that interpret the more general norms in an international convention. These rules are mostly intended to add world knowledge that may otherwise stay implicit. For instance a rule that the capacity of fire extinguishing pumps ought not to exceed the capacity of the bilge ejection system by some specified margin, and another rule that it is prohibited to transport flammable liquids through a piping system that would be used to transport a fire extinguishing medium in case of fire, are both affirmations of the more general rule that safety arrangements prone to human error with possibly disastrous consequences are not allowed. This relation is somewhat similar to that between a norm and underlying principles and/or goals (cf. [8]). It is also the obvious dual of disaffirmation. This relation is not used directly in argumentation, but it is very convenient to find norms by general purpose or topic.

Disaffirmation – A norm disaffirms another one if it regulates a more specific case, but is classified in opposition. This definition is in line with Valente in [26]. Disaffirmation is defined by Lindahl (cf. [10]) as “A conflict between two norms of different deontic mode, one being permissive and the other mandatory”. This is the common type of exception between norms and this way of representing them has been explored in detail in i.a. [30]. Usually the more specific norm, the one that disaffirms, has been added with the specific intent of creating an exception for a (more) specific situation. This works if *Lex Specialis* is the strongest principle applicable for resolving the conflict, and the disaffirmation is asymmetric. If the latter is not true, we have a case of disaffirmation conflict. Disaffirmation is used to construct rebutting arguments (as in the example).

Disaffirmation Conflict – A relation of disaffirmation conflict is a symmetric disaffirmation that cannot be resolved through the *lex specialis* principle. This situation is absurd, and likely the consequence of a modelling error (usually in the world theory) or a

legislation error. This case is incorrect and must be immediately resolved by the indexer that introduced it.

Compliance Conflict – What if the cases to which two norms apply cannot be jointly realised? If at least one of them would have been a permission, and one of them is applicable to the case at hand, then it would have been futile to try to prove that the case is a theorem of the case classified by the other norm. The case is not regulated. If both norms are classified as disallowed, we have case of compliance conflict (cf. [10]). Compliance conflicts are permissible, but always reported to the client, whenever one of both involved norms is applied in dialog by the plaintiff. Note that because the proof of disjointness is by definition symmetric, a compliance conflict cannot be resolved through *Lex Specialis*, unless some preference criterium for proof construction is included. Valente uses a special notion of negation (opposition) in [26] to make such a distinction.

The disaffirmation link is used to construct and warrant rebutting attacks. The affirmation link allows an indirect rebuttal: First try to apply an affirming norm, and if it does, apply the norm that disaffirms it. This strategy is more rigorous than the winning strategy presented before, but indirect rebuttals appear to change the topic of the dialog.

Including document elements in IR results that disaffirm a selected document element or create a compliance conflict with it always increase recall. The exact effect depends on the question: less facts is better. Precision remains mostly a function of the accuracy of the match procedure between question and document elements – the expressiveness and comprehensiveness of the theory about the regulated world in the game. Document elements linked up with affirmation and disaffirmation links produce a tree (forest, more accurately) whose branches can be methodically explored, starting at the root(s), and be cut off if a document element is rejected by the user. It is possible to compute the tree from case descriptions and the domain theory, but it may be far more efficient to use a reversed Consultancy Game (where the system plays the client using a knowledge elicitation if there are no real clients available to interface to the expert) with legal experts to construct the tree, the cases, and the concepts in the domain theory through argumentation.

7. Discussion

Results of forcing some actual transcripts of legal advice sessions into the mould of the Consultancy Game are encouraging. Rules are typically referenced in the order prescribed by the game, even though the conversation is less structured. In contrast, discovering the links between document elements from their content is far more difficult.

There are two subject representation problems that require further discussion. Firstly, scoping is a problematic concept. A rule that applies to a case, but has no normative impact, can be a scoping rule. This concept is important because it solves a representation problem that will occur in any legal document. It also states that this case to which the scoping rule applies subsumes all cases to which any rule indicated in the scope applies – the scoping rule itself describes a relation over rules. It describes another domain – that of the structure of the document – and is semantically grounded in that domain. This is more elegant than trying to mix this information into the knowledge representation describing what the document is about. This distinction can however get very fuzzy; Let's take an example that is classified by Prakken and Schrickx in [20] as an implicit exception; Section 1624 of the Dutch Civil Code (BW) declares that if a rule about rent contracts for business accommodation conflicts with a rule about another type of contract, the former rule is prior to the latter. It is clearly necessary to read the rules, at least the introductory sections where one will find scoping rules, to construct the priority ordering for a rule like this. This example affects the choice for priority ordering on norms. The ordering is actually on

norms instead of document references because a rule may have multiple positions in the ordering as it applies to different cases. In previous example this case arises if a rule about another type of contract applies to multiple cases, only some of which are consistent with rent contracts for business accommodation.

Secondly, the formalisation of compliance conflict as intended by Lindahl [10] is not entirely satisfactory. The relevance of the conflict is largely determined by who will be held responsible for violation for any of the norms involved. If an agent is prevented from complying with both norms, the agent is not a priori responsible for the resulting violation of one of the norms. The agent may not perceive any compliance conflict. Responsibility issues are less suitable for indexing in general, because argumentation about responsibility requires very complex reasoning, even if it rarely departs from a common-sense assignment of responsibility. This type of reasoning does not lend itself to high precision IR, and there is probably little need for it in a consultancy context, since rules about establishing responsibility are typically few in number.

References

[1] Ambroziak, J., & Woods, W. (1998). Natural Language Technology in precision content retrieval. In: Proceedings of the International Conference on Natural Language Processing and Industrial Applications. Moncton, New Brunswick, Canada, 1998.

[2] Blair, D.C. and M.E. Maron (1985). An evaluation of retrieval effectiveness for a full-text document retrieval system. Communications of the ACM 28(3):289-299.

[3] Breuker, J.A. (1992). On Legal Information Serving. In: Proceedings of JURIX'92, Information Technology and Law. Koninklijke Vermande, Lelystad, pages 93-102.

[4] Breuker, J. A. & N. den Haan (1991). Separating regulation from world knowledge: where is the logic?, in M. Sergot (Ed.), Proceedings of the 4th International Conference on AI and Law, pages 41-51 (New York, ACM).

[5] Breuker, J. A., Valente, A. (1994). Ontologies: the Missing Link between Legal Theory and AI & Law. In: Proceedings of JURIX'94, pp. 139-150. Vermande, Lelystad.

[6] Greenleaf, G., Mowbray, A., & Lewis, D. (1994). Australasian Computerised Legal Information Handbook. Butterworths, Australia, 1994.

[7] Hafner, C.D. (1981). An Information Retrieval System Based on a Computer Model of Legal Knowledge. PhD thesis, University of Michigan, 1981. Republished by UMI Research Press, Ann Arbor, Michigan, 1981.

[8] Hage, J. (1995). Teleological Reasoning in Reason-Based Logic. In: Proceedings of ICAIL-99, ACM, New York, pp. 11-20, 1995.

[9] Hamblin, C. (1971). Mathematical Models of Dialogue. Theoria, **37**, 130-155.

[10] Lindahl, L. (1992). Conflicts in systems of legal norms: A logical point of view. In P. Brouwer, Bt. Hol, A. Soeteman, W. van der Velde and A. de Wild (eds.). Coherence and Conflict in Law, pp. 39-64. Kluwer, Deventer, 1992.

[11] Lodder, A. (1998). On Structure and Naturalness in Dialogical Models of Argumentation. In: Proceedings of JURIX-98, pp. 45-58. GNI, Nijmegen.

[12] Miller, G. A., Beckwith, R., Fellbaum, C., Gross, D., and Miller, K. J. (1990). Introduction to WordNet: an on-line lexical database. In: International Journal of Lexicography 3 (4):235 – 244, 1990.

[13] Nohr, H. (1999). Inhaltsanalyse. NFD Information – Wissenschaft und Praxis, 50(2):69-78.

[14] Palipoli, L., Pirri, F., Pizzuti, C. (1999). Algorithms for selective enumeration of prime implicants. Artificial Intelligence, **111**:41-72 (1999).

[15] Van der Pol, R. W. (2000). Knowledge-based Query formulation in Information Retrieval. Phd thesis, University of Maastricht, 2000.

[16] Pollock, J.L. (1987). Defeasible Reasoning. Cognitive Science **11**:481-518.

[17] Power, R., D. Scott, and R. Evans (1997). What You See Is What You Meant: direct knowledge editing with natural language feedback, Technical Report, ITRI-97-03, University of Brighton, 1997.

[18] Prakken, H. (1993). Logical Tools for Modeling Legal Argument. PhD Thesis, Free University of Amsterdam, 1993.

[19] Prakken, H. and Sartor, G. (1995). On the relation between legal language and legal argument: assumptions, applicability and dynamic priorities. In: Proceedings of the fifth ICAIL, ACM Press, Washington DC, pp. 1-9, 1995.
[20] Prakken, H. & Schrickx, J. (1991). Isomorphic models for rules and exceptions in legislation. In J.A.P.J. Breuker, R.V. de Mulder, J.C. Hage(eds.): Legal Knowledge Based Systems. Model-Based Legal Reasoning (Proceedings of JURIX'91), 17-27. Koninklijke Vermande BV, Lelystad, 1991.
[21] Schlag, P. (1996). Hiding the Ball. New York University Law Review **1681**, 1996.
[22] Sigel, A. (2000). How can user-oriented depth analysis be constructively guided?. In: Beghtol, Clare; Howarth, Lynne C. & Williamson, Nancy J. (eds.). Dynamism and Stability in Knowledge Organization. Ergon Verlag, Würzburg, 2000.
[23] Straach, J., Truemper, K. (1999). Learning to ask relevant questions. Artificial Intelligence, **111**:301-328 (1999).
[24] Turtle, H. (1995). Text retrieval in the legal world. Artificial Intelligence and Law, 3(1-2):5-54, 1995.
[25] Uyttendaele, C., Moens, M. F., DuMortier, J. (1996). Salomon: automatic abstracting of legal cases for effective access to court decisions. In: Proceedings of JURIX'96, Tilburg University Press, 1996.
[26] Valente, A. (1995). Legal Knowledge Engineering. A Modeling Approach. PhD Thesis, University of Amsterdam. IOS Press, Amsterdam, 1995.
[27] Winkels, R.G.F. (1998). CLIME: Legal Information Serving Put to the Test. In: Pre-proceedings of the Second French-American Conference on AI and Law.
[28] Winkels, R., De Bruijn, H. (1996). Case Frames: Bridging the Gap between a Case and the Law. Proceedings of the First European Conference on Computers, Law and AI, pp. 205-213, 1996.
[29] Winkels, R.G.F., Boer, A., Breuker, J.A. and Bosscher, D. (1998). Assessment Based Legal Information Serving and Cooperative Dialogue in CLIME. In: Proceedings of JURIX-98, pp. 131-146. GNI, Nijmegen, 1998.
[30] Winkels, R.G.F., Boer, A., Breuker, J.A. and Bosscher, D. (1999). Generating Exception Structures for Legal Information Serving. In: Proceedings of ICAIL-99, ACM, New York, pp. 182-189, 1999.