

Case Reconstruction before Assessment

Henk de Bruijn and Radboud Winkels

Department of Computer Science & Law

Faculty of Law, University of Amsterdam

email: {henk,winkels}@lri.jur.uva.nl

Abstract

Traditionally Legal Knowledge Based System (LKBS) research concentrates on the matching of abstracted case facts with some formal representation of the law. How these facts are obtained, and how the abstraction that is necessary to 'match' the facts with terms mentioned in legal rules or precedents, is accomplished, is left unmentioned (typically the user is assigned this task). The paper will describe this 'reconstruction' phase that precedes legal assessment of the case at hand. An elaboration of exactly this subtask can solve many problems of developing LKBS's.

The development of a practical application for first order legal aid to members of one of the largest trade unions in The Netherlands by non professional volunteers is described. A LKBS is being developed to improve referral to professionals and to induce a better understanding of the law by the volunteers when handling 'dismissal' cases. Analysis of the tasks executed by the volunteers reveals the necessity of extending the core of legal reasoning, assessment, with a practical part, aimed at supporting the reconstruction of the case. The support is based on stereotypical plans (scenario's), and an interface that reflects these scenario's. Furthermore, the application is embedded in a traditional information retrieval environment. A successful prototype has been developed, and we are now in the phase of building the actual system.

Keywords: assessment, legal case reconstruction, stereotypical plans, real-world application

1 Introduction

This paper describes a research project which aims at designing and implementing a practical application using artificial intelligence techniques in a legal domain: a Legal Knowledge Based System (LKBS) for Dutch dismissal law. The notion that the implemented system is going to be used in practice is important. As a result, the research is focussed on the more practical problem of reconstruction of legal cases, instead of a new approach of matching facts and rules. Legal information and knowledge-based systems can

only function on the basis of a *complete and coherent case description*. Present systems do not offer any substantial support in obtaining such a description. Typically the user (legal professional or layman) has a limited set of terms at his/her disposal, and describing the case at hand consists of 'selecting' the appropriate terms, usually prompted by (multiple-choice) questions from the system. The user has no overview over the case facts provided, and relations between facts are absent or implicit. The resulting case description may be incomplete, and will most likely be incoherent. In a coherent case the different roles of objects and subjects in situations/actions (*identity*), and the causal and teleological relations between situations/actions (*causality and intentions*) are clear. When these are not clear, the system will not be able to reason about motives and responsibilities, nor about possible future actions, and only partially about the applicability of legal terms and norms.

In the research described here, we focus on supporting the intake of full case descriptions. We suggest the use of stereotypical, abstract case frames (plans), and a transparent user interface that reflects these plans at a conceptual level that compares to the level of understanding of the user. The case frames used, represent the coherence of cases more explicitly, when roles of agents and actions in time are concerned. This is the reason for calling these case frames 'skeletal plans'. Using these plans facilitates a more natural style of presentation to the user. In order to allow also for a more flexible style of reasoning about the coherence of plans, it is suggested to investigate an even more profound representation of planning knowledge.

Another way of supporting users with the abstraction of legal cases is introduced by using an intermediate level of conceptual knowledge in the collection of prototypical cases. The terminology used for constructing and presenting these cases stems from (a.o sources) educational material, instead of legal sources. Legal definition rules and compiled out expert knowledge is used to obtain a case description at the right level of abstraction.

The research started with building a model of expertise for legal consulting. The model of expertise stems from the CommonKADS methodology, a standard in knowledge engineering [Schreiber et.al 1993]. A model of expertise consists of three layers. The first layer consists of a task decomposition and a control structure for the process of reasoning. The second layer encompasses the necessary inferences: reasoning steps capable of inferring new knowledge, e.g. abstraction of a case, or selection of a stereotypical plan (see below). Finally the third layer contains the domain knowledge that needs to be represented in order to make reasoning

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

possible. Representations of legal rules belong to this last layer.

The description of the research starts with an outline of the motive for and the purposes of the project. The development of an expertise model and design of the implementation is reported upon after this. In our case, this phase is based upon an analysis of the functions that have to be fulfilled inside the organization. Empirical and analytical research of the task of (legal) consulting reveals the necessary subdivision of the global task. Subtasks as collecting and abstracting facts and using the resulting abstracted case description to solve the present problem will be identified. We intend to support these various subtasks with the implemented system in different ways. After the evaluation of the prototype, a model for legal consulting is suggested. An example of the model, with relevant aspects of the domain is given. The paper is concluded by the lessons learned from this practical research.

2 Motive and purposes

The system is being developed for one of the largest trade unions in Holland, the Industrial Union FNV. This union has about 250.000 members. In principal, first legal aid for the benefit of its members is provided by volunteers. In the near future, these volunteers will be concentrated in approximately eighty law centres, spread over the country. This development consolidates the deployment of volunteers in order to take over tasks previously performed by professionals. The research project described in this paper, started with an investigation of the possible negative consequences of this development and the possibilities for (knowledge based) automation of some tasks at the law centres. Problems have been identified concerning different aspects of the quality of the offered services: consistency, and reliability of the advice given, and the first impression the clients get. The volunteers tend to express a lack of knowledge, resulting in insecurity and a non-efficient way of problem solving. As a result, a great number of case files is still referred to legal professionals, lacking the necessary details.

Several solutions for the problems mentioned exist: education of the volunteers, support by means of hardcopy documentation, and the use of computers for assistance or partial automation of the tasks to perform. The union is already investing a great deal of time and money in education. The problem with existing documentation is the difficulty in updating it, and smooth coupling with the knowledge level of the volunteers. The last, uptill now underdeveloped alternative, the use of (knowledge based) information technology, seems to be the best opportunity to improve the quality of legal aid to the members of the union.

The system to be implemented will serve two main goals. First the referral of cases to professional lawyers should be improved by obtaining all data from the client, necessary for the solution of the presented problem. In other words, a complete and coherent case description should be obtained. Secondly, the use of the system should provoke a better understanding of the case by the volunteer. This will lead to a more satisfying settlement of the case from the perspective of the client.

It is important to note that deciding a case is not the main goal for the described system. Collection of essential data and pointing into possible directions is. Of course, in order to decide a case, these subtasks are crucial.

Expert: '...Well, I assume A. has a labourcontract, which implies a procedure for an appeal for dismissal.

....

Furthermore, it is important to know to what extend A. can be forced to work at regular times. Did she agree on terms of restricted availability?'

Volunteer: '....She has agreed to call every week to ask where and how she might work. So it is rather easy to put her aside'

Expert: 'It might be not so bad. It depends on the question whether she is obligated to appear on demand....'

Figure 1: An example of the results of the Wizard of Oz experiments

3 Task analysis

The analysis of the consulting task aimed at the way in which this task should be executed. Also, the analysis addressed how the future users currently perform this task in practice. Different methods were used to carry out this analysis. In the first place an extensive study of educational material and directions of the volunteers was combined with a rational task analysis [De Greef & Breuker, 1992].

Secondly, the resulting task decomposition was revised using the results of a so called 'Wizard of Oz' experiment [Winkels, 1992]. During such an experiment, an expert is acting as a mock up of the prospective system, which is consulted by a future user. The texts exchanged during interaction are recorded just like the voices of both expert and user who are asked to think aloud. In our case, volunteers were asked to solve a dismissal case, with support of an expert. Volunteer and expert worked in different rooms during the experiments and interacted via terminals in natural language. In this way data have been collected concerning task conception, knowledge level and expectations of the user. Moreover, these data facilitated the comparison of the suggested task decomposition with the actual behaviour of experts.

The Wizard of Oz experiments lead to the following observations. At a very abstract level, the terminological knowledge of the volunteers was sufficient in order to communicate with the expert. At the other hand, the incomplete, and sometimes erroneous knowledge of the volunteers lead them to jump to conclusions; typical novice behaviour (e.g. [Jansweijer, 1988]). The experts used the terminology and interpretations of the volunteers, whenever possible, but also tried to 'broaden their horizon' by stating alternative interpretations, and asking for additional information. Our mock-up experiments once more stressed the importance of involving the consulting person (in this case the volunteer) in the problem solving process, for example by giving an opportunity for generating and specifying hypotheses. The task analysis resulted in a task decomposition, consisting of subtasks like listening, formulating the clients problem, information retrieval etc. These subtasks can be divided into three functional categories: collecting and specifying the data, combining this data with existing information and knowledge, and finally producing a referral to the professional lawyer. Our observations correspond to earlier findings on consulting [Johnson et.al, 1987]. In this research, the way an (artificial) consultant should perform his task

was considered. Abstraction of the collected data into the terminology of the domain of expertise, and refinement of the data in cooperation with the client is emphasized.

4 Task distribution in the prototype

Distribution of the various subtasks over the involved agents led to the outline for the architecture of the prototype system. The involved agents are: volunteer, LKBS, and the already existing database UIS (Union Information System). In principle, the user is in full control of the task execution. Support is given by two knowledge based modules: a rule based assessment module and a reconstruction module that is rooted in a skeletal frame for dismissal cases. The complete knowledge based system uses one data representation, a frame based model of the world: concepts, attributes and actions. The UIS supplies data about members which are already recorded. In addition, the UIS includes a text retrieval system. Texts originating from educational material and supporting documentation can be withdrawn by the user with key words.

When a union volunteer is confronted with someone who is (possibly going to be) fired, the session begins with the retrieval of data already existing in the UIS. The interaction between the LKBS and the user follows next. Since the user is in control, he or she can decide which information to supply at what time. The system does not force the user in a specific sequence of questions. However, the case description phase is prestructured by using a skeletal frame for dismissal cases, consisting of four phases: entering into the labour contract, the causes for dismissal, the actual dismissal and any counteraction after dismissal.

These four different stages have been identified during the domain analysis. It is very difficult to reason about cases considering agreements, without some notion of successive actions (cf. [Gardner, 1987]). The frame is presented to the user, using different screens (forms) for all stages. The layout of an example screen is presented in Fig.2.¹ The user is given a complete overview of possible relevant facts, concerning the selected phase. Furthermore, jumping to other phases and options is possible at all times.

A central role for the volunteer handling the case is stressed. Apart from assigning the main part of the control of the global task and the execution of the abstraction task to the user, he or she is also able to pose hypotheses concerning the case, that can be matched with the legal rules.

5 Evaluation of the prototype

Users evaluate the use of forms, representing stages corresponding to their natural conception of 'stories', positively. The same holds for the aspect of user control. However, an important problem remains.

The abstractions of volunteers, concerning case facts, do not correspond to legal terminology and cannot directly be matched against legal rules. The volunteers use a strategy to simplify the complexity of handling legal cases, that seems to be based on the recognition of prototypical cases. In this way, the central issue of a case can be spotted quickly, which makes gathering relevant facts and recollecting relevant knowledge easier. The conclusion that prototypical cases play an important role in making legal cases acces-

¹An English version is presented. The original interface is in Dutch.

SCREEN	ACTUAL DISMISSAL	OPTIONS
Type of dismissal	Termination notice Appeal for dismissal Rescission magistrate Instant dismissal End of contract Probationary period	Data employee Data employer Facts about contract Motive for dismissal Actual dismissal Procedure so far
Date of dismissal	25.05.1995	Hypothesis Evaluation Referral
Conveyance	Written notice Verbal notice	Help Quit

Figure 2: An example of the screens representing successive stages in dismissal cases. The user is able to manipulate data s/he knows about, or to jump to another phase, considered to be more relevant.

sible for beginners is supported by experiences during the experiments mentioned above.

The idea of using prototypical cases, or case frames in order to facilitate selection and abstraction of relevant facts is elaborated in the next sections. These sections start with a model which plays a central role in most research of LKBS's: a model of assessment.

6 Towards a model of legal consulting

Assessment is the core of all legal reasoning, as demonstrated by e.g. Breuker [Breuker, 1992]. Assessment is essential for all kinds of legal reasoning, including reasoning on behalf of retrieval of legal texts. The main point of assessment is the match of a normative system against an abstract case description [Breuker et. al, 1987].

The mere remark that such a task is central to legal reasoning will not be surprising. Most existing legal expert systems mainly depend on matching formalized legal rules with collected facts. However, a direct matching procedure, driven by successive one-sided questions and answers, is bound to come to a deadlock.

New insights in solving the problem of preparing the match in assessment have been presented in recent research in the field of assessment [Valente & Löckenhof, 1994]. Descriptions are given for different configurations of inference steps, covering the *abstraction* of case descriptions and the *match* of these abstractions with the normative system.

First, the translation of accumulated facts into legal terminology can happen in two ways: two-step abstraction and using case frames. In case of a two-step abstraction, unstructured data is structured and selected in order to acquire relevant case data. This can be done by form filling for instance. When someone is confronted with a form, while entering data about a case at hand, s/he is forced to adapt to the (implicitly) represented coherence of the case. If the relevance (or coherence) of data differs for different cases, case frames can be used. When the case at hand is recognized as 'typical', the exact case frame focusses on relevant data with a predisposed coherence.

It is easy to notice that in the prototype we used a two-step abstraction. Relevant case data were extracted by filling four different forms for four different stages. Although the use of different stages implies some sort of skeletal plan or case frame, it is important to consider the fact that no difference is made concerning the relevance of facts, as this is exactly the main goal of the use of case frames. In the next section we will propose a model for legal consulting which *does* differentiate between the relevance of facts for different typical cases.

The inferences described above, aim at another important step: abstraction of the facts, as well. In LKBS's, abstraction can be guided by the use of definitional legal knowledge. Many legal rules are concerned with defining legal concepts. In the proposed model for legal consulting, these kind of legal rules will play an important role in arriving at an abstract case description. After completion of the abstraction inference, it should be possible to match the case description with a selected set of rules.

The selection of appropriate rules in order to perform the necessary *match* is the second branch developed in the model of assessment, although not a key issue in this paper. This selection also includes the recognition of conflicts between rules and the attachment of priority values. A justified appreciation of these features is essential for reaching a right formalization of legal rules.

In our view the development of these two branches (abstract and match) coincide with the proposed separation of world (*terminological*) knowledge and legal (*normative*) knowledge in legal knowledge based systems by Breuker and Den Haan (e.g. [Breuker & den Haan, 1991]). They propose such a separation in order to reach a tractable matching procedure. If matching is tractable, the use of conditional answers and hypothetical reasoning is facilitated.

These reflections result in a model for legal reasoning proposed by Valente and Breuker [Valente & Breuker, 1994] and presented in Fig. 3. Structuring the facts and reasoning about terminological (world) knowledge is separated from the matching of the abstracted case-description with formalized rules. Further on in this paper, an elaboration of reasoning with world knowledge will be presented.

7 A model of reconstruction in legal consulting

The system now under development, aims at the reconciliation of lay knowledge, expert knowledge and (written) legal knowledge. Lay knowledge is used for the specification of facts and the recognition of skeletal (stereotypical) plans. The expert knowledge is incorporated in the composition of the plans and compiled out in a small part of surface knowledge to facilitate abstraction. Legal knowledge is used for abstraction (definitional knowledge) and assessment (normative knowledge).

A part of the proposed inference structure for legal consulting that is used for designing the architecture of the system is presented in Fig. 4. The figure represents the inference steps and domain knowledge, necessary for executing the abstraction task.

After the collection of some initial facts, very close to a (for the user) natural conception of events in the world, a first abstraction can be made. This abstraction consists of the selection of a skeletal plan. In a skeletal plan data about essential facts are represented. These data are used during specification of the skeletal plan, by determining the essential facts. The way these skeletal plans (or case frames) and the terminology that is used in these plans, compares

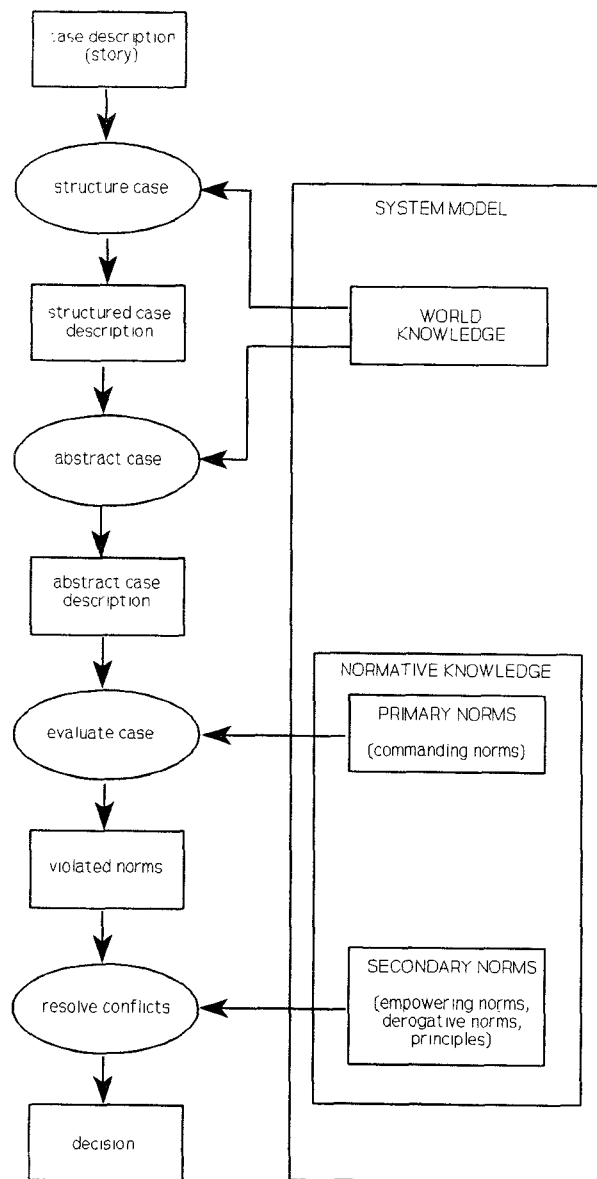


Figure 3: Model for legal assessment, developed by Valente & Breuker, 1994.

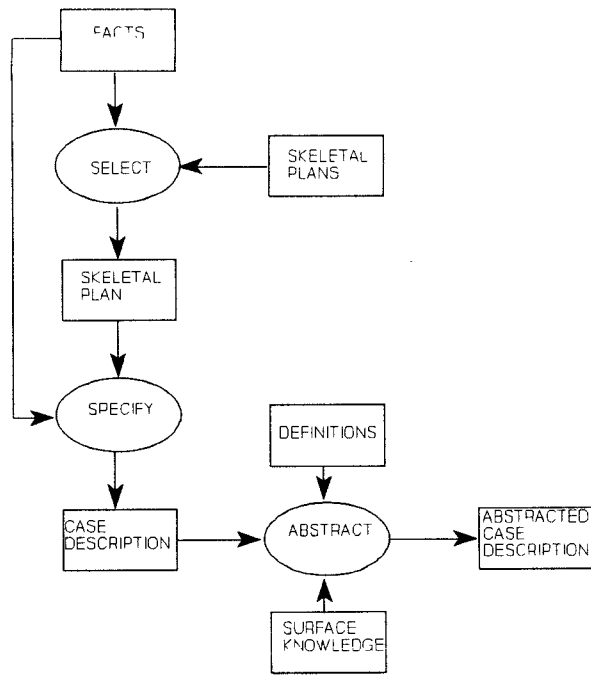


Figure 4: The proposed inference model for collecting and abstracting relevant facts, precedent to matching facts and legal rules.

to legal terminology, is explained in a detailed example in the next section.

After the first steps of abstraction (structuring the data) the final step towards an abstracted case description in legal terminology can be made. Two types of knowledge are used to perform this final abstraction step: legal (definitional) knowledge and compiled out expert knowledge (surface knowledge). The resulting abstracted case description can be matched against the formalized rules.²

In the next section an example of the use of domain knowledge in this inference model is presented.

8 An example: dismissal of on-call workers

Special contracts for on-call workers³ are not defined in Dutch labour law. Only two sorts of labour contracts are defined: contracts for definite and indefinite duration. It is not possible to deduce straight forward from the rules in the law, whether an on-call worker is in the possession of a labour contract or not. Three conditions must be met before the conclusion that a labour contract exists can be drawn: the relation between employer and employee should be distinguished by authority of the employer, labour has to be performed by the employee personally and he should be paid for this. The role these conditions play in the factual decisions of judges is represented in a number of precedent

²Currently, this approach, using a weak problem solving model precedent to a strong model is an accepted way of combining the known advantages of respectively swiftness and accuracy of such models [Simmons, 1992].

³On-call workers did not agree for a fixed number of working hours. The employer can ask them to work when necessary.

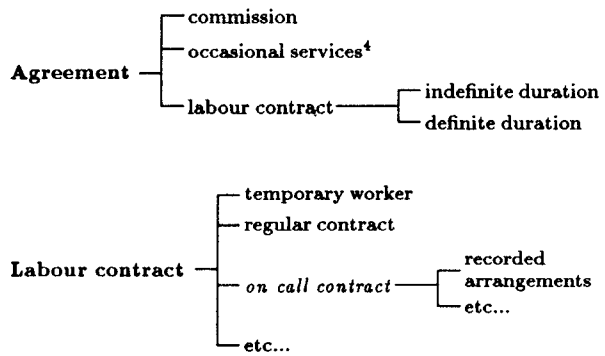


Figure 5: Two hierarchies concerning contracts in dismissal law, the hierarchy above represents the possibilities mentioned in the law, the other hierarchy gives an partial overview of known descriptions of real world situations

cases [Bakels, 1990; Bakels, 1992]. The feature of authority plays the most central role in deciding these cases. So, it is not the agreement two parties have signed on paper that is at stake, but the actual content of the relationship between employer and employee. The issue of labour contracts is very important in dismissal cases, because a far reaching employment protection exists, which in principle forbids dismissal if such a contract (or an equivalent relationship) exists.

The result of this situation is that the volunteers of the union are confronted with a great number of cases of on-call workers that appear to be the same, but will result in very different outcomes, depending on the existence of a labour contract. Fortunately, experts are able to give some rules of thumb, which reflect the way judges tend to decide these sort of cases. Important issues are i.a.: the continuity of the relationship between employer and on-call worker, the party who initiated arrangements for jobs, and possible problems for the employee if he resisted to agree with such initiatives. All these facts contribute to deciding whether authority is present or not. However, none of these facts are represented in the law straight away. A skeletal plan for cases concerning on-call workers focusing on these central issues resembles the knowledge embedded in precedent cases as well as in legal experts. Specifying the indicated data should be the starting point for solving such cases.

The way the notion of on-call workers compares to terminology used in labour law, is roughly represented in figure 5. The problem is that the two represented hierarchies do *not* match. The concepts in the different hierarchies belong to a different abstraction level. For every concept existing in the real world, it is important to determine the position in the legal terminology. We claim specific abstraction rules for the distinct situations (and concepts) in typical cases exist. In case of on-call workers for example, some of these rules are given in the example below.

Note that the most relevant facts in cases concerning on-call workers, are concentrated in the first phase: entering into the labour contract. Using case frames in stead of only four phases to concentrate on relevant facts, facilitates a more specific grouping of relevant facts in the case frame that is chosen.

Once the existence of a labour contract is demonstrated, the evaluation of the actual dismissal can begin. Different ways of firing people are defined, like rescission by a magis-

trate, instant dismissal, appeal for permission for dismissal, etc. Each form of dismissal has its own possible motivations and moreover specific procedural rules to be followed.

The way the system under implementation will support the intake of cases described above will be clarified in the following specification of the named inference steps.

Selection of the skeletal plan: The first step consists of deciding which plan is going to be used during specification of the facts. A list of possible descriptions of the present case is presented to the user. The terminology used originates from practice and suits the way the volunteers are educated, so the selection of the proper items should be no problem for the user. This terminology compares to the hierarchy at the bottom of figure 5. Global terms and specific attributes selected by the volunteer are compared by the system in order to be able to select the prototypical plan.

Concerning on-call workers: treating on-call workers as a category is very familiar to the volunteers. Answering the question whether arrangements have been recorded on paper leads to an even more specific case frame. The chosen case frame facilitates the next step.

Specification of the facts: The attributes pointed out by the expert as being essential for on-call workers, concerning authority, are presented to the user in shape of sequential forms. These forms again represent the different stages of a dismissal case. The user is now in the position to specify the facts of the present case.

If both parties have recorded their arrangements, aspects with increased relevance are for example: the party having the initiative for entering new engagements, a possible obligation for the employee to conform to the wishes of the employer etc. If recorded arrangements do not exist, comparable aspects are relevant, but cannot be checked directly. In this case, relevant facts like past conflicts about vacations (in connection with obligations) and the existence of fixed working-days (supported with pay slips for example) have been included in the case frame.

Abstraction: The system combines the definitional knowledge from the law with the expert knowledge about precedent cases. An abstraction at level of legal terms (the hierarchy at the top of figure 5) should be reached. In figure 6 two examples of abstraction rules are given. When assessing a complete case, not only the existence of a labour contract is determined (abstracted). For example all the actions of the employer in connection with the dismissal are abstracted in order to match the procedural rules.

Match: The resulting abstracted case description is matched to formalized rules. An example, describing such a match can be found in [Den Haan & Breuker, 1991].

9 Conclusion

We take the position that in order to build a practical LKBS, one should analyse the task that has to be executed, in stead of taking a fixed model of legal reasoning for granted. Concerning the task of legal consulting, the task analysis presented in this paper leads to the conclusion that gathering and abstracting relevant facts is at least as important as matching these facts with the rules.

If a contract includes:
Personal labour &
Payment &
Authority in relationship between parties
Then
This contract is a labour contract

If An on call contract includes:
An obligation to appear on demand
Then
This contract includes authority in relationship between parties

Figure 6: Examples of abstraction rules in law and expert knowledge respectively

It is a mistake to place the main part of the burden of abstraction on the users of LKBS's. Although people outperform computer systems on reasoning about the world, efforts trying to support such tasks should be encouraged. Two ways of doing this are presented in the paper.

Firstly, we supported reconstruction with an adequate interface, adjusted to the (task) conception of the user. Users evaluate the presentation of forms for different phases of dismissal cases positively. This interface replaces a sequence of machine driven answers and questions.

Secondly, compiled out expert knowledge can be used for configuring a set of skeletal plans, which comprise knowledge about the essence of certain facts in recognized typical cases. Furthermore, knowledge about the way judges used essential facts in precedent cases, should be represented and reasoned with, in order to support abstraction of facts to legal terminology. Once an abstracted case description is reached, the computer should do what it is good at: reasoning with rules.

When trying to support abstraction of legal cases, one should recognize the fact that humans do not make abstractions to legal terminology in one step. Different abstractions are used (like in text books) to improve accessibility of complex domains. In combination with skeletal plans, these abstractions are useful for identifying relevant facts. This was demonstrated with an example about classification of labour contracts.

Supporting abstraction should, to some extent, incorporate world knowledge. From this point of view, it might be promising to improve reasoning about reconstruction (and planning) in LKBS's, with an even more elaborate use of planning techniques, surpassing the use of skeletal plans. Although the use of skeletal plans implies the choice for an explicit coherence between the represented facts, maybe this coherence should be represented more explicitly, using elements of planning knowledge. As we are modelling a domain heavily involved with concepts stemming from the planning domain, like agents, actions, intentions, goals and assessment, this approach appears both natural and promising. After all, we are modelling a domain heavily involved with concepts stemming from the planning domain, like agents, actions, intentions, goals and assessment. Recent research, using planning knowledge in the representation of a legal domain, supplies arguments in support of this view [Visser et.al, 1994]. This research concentrated on planning of actions, in stead of reconstruction. The difference may be surprisingly small though.

References

- H.L. Bakels. *Arbeidsrechtspraak; civiele rechtspraak in arbeidszaken geselecteerd en geannoteerd door Prof.mr. H.L. Bakels*. Kluwer, Deventer, 1990.
- H.L. Bakels. *Schets van het Nederlands arbeidsrecht*. Kluwer, Deventer, 1992.
- J.A.P.J. Breuker. On Legal Information Serving. In *C.A.F.M. Grütters, J.A.P.J. Breuker, H.J. van den Herik, A.H.J. Schmidt and C.N.J. de Vey Mestdagh (eds), Legal Knowledge Based Systems: Information Technology and Law, JURIX'92*, Koninklijke Vermande, Lelystad, NL, 1992.
- J.A. Breuker & N. den Haan. Separating World and Regulation Knowledge: Where is the Logic? In *Proceedings of the Third International Conference on Artificial Intelligence and Law*, Oxford, 1991.
- J.A. Breuker, B. Wielinga, M. van Someren, R. de Hoog, G. Schreiber, P. de Greef, B. Bredeweg, J. Wielemaler, J-P. Billault, M. Davoodi, & S. Hayward *Model Driven Knowledge Acquisition: Interpretation Models*. ESPRIT Project P1098 Deliverable D1 (Task A1), University of Amsterdam and STL Ltd, 1987.
- P. de Greef, & J.A. Breuker. Analysing System-User Cooperation in KADS. In *Knowledge Acquisition*, 4(1), pp.89-108, 1992.
- N. den Haan & J.A. Breuker. A Tractable Juridical KBS for Teaching and Applying Traffic Rules. In *JURIX 1991 - Fourth International Conference on Legal Knowledge Based Systems*, Koninklijke Vermande, Lelystad, NL, 1991.
- A. von der Lieth Gardner. *An Artificial Intelligence Approach to Legal Reasoning* MIT Press, Cambridge, 1987.
- W. N. H. Jansweijer. *PDP: An artificial Intelligence Approach to Problem Solving and Learning by Doing in a Semantically Rich Domain*. PhD thesis, University of Amsterdam, Amsterdam, The Netherlands, 1988. (in Dutch).
- P.E. Johnson, C.J. Nachtshein & I.A. Zualkernan. Consultant Expertise. In *Expert Systems*, 4(3), pp. 180-188, 1987.
- G.B. Schreiber, B. Wielinga, J.A. Breuker (eds). *KADS: a Principled Approach to Knowledge Based System Development*. Academic Press, New York, 1993.
- R.G. Simmons. The Roles of Associational and Causal Reasoning in Problem Solving. In *Artificial Intelligence*, 53, pp.159-207, 1992.
- A. Valente & J.A. Breuker. Making Ends Meet: Conceptual Models and Ontologies in Legal Problem Solving. In *Proceedings of the 11th Brazilian AI Conference, SBIA'94*, 1994.
- A. Valente & C. Löckenhof. Assessment. In *J. Breuker and W. van de Velde, eds. Reusable Problem Solving Components: The CommonKADS Library for Expertise Modelling*, IOS Press, Amsterdam, 1994.
- P.R.S. Visser, R.W. Van Kralingen & H.J van den Herik. Towards a Control-independent Representation of Legal Knowledge. In *Proceedings of the Fourth National Conference on Law, Computers and Artificial Intelligence*, University of Exeter, pp. 166-176, 1994.
- R. Winkels. *Explorations in Intelligent Tutoring and Help*. IOS Press, Amsterdam, 1992.

Acknowledgement

Ideas expressed in this paper also evolved in discussions with Joost Breuker. We would also like to thank the anonymous experts and volunteers for their participation in the experiments. The Industrial Union FNV contributed both people with good ideas and funds. Finally, we thank the anonymous reviewers for their comments on the earlier version of this paper.