Legal knowledge based systems
JURIX 92
Information Technology and Law

The Foundation for Legal Knowledge Systems
Editors:
C.A.F.M. Grütters
J.A.P.J. Breuker

H.J. Van den Herik
A.H.J. Schmidt
C.N.J. De Vey Mestdagh

P. Mariani, D. Tiscornia and F. Turchi, The Formalization of Retrieval and Advisory systems, in: C.A.F.M. Grütters, J.A.P.J. Breuker, H.J. Van den Herik, A.H.J. Schmidt, C.N.J. De Vey Mestdagh (eds.), Legal knowledge based systems JURIX 92: Information Technology and Law , The Foundation for Legal Knowledge Systems, Lelystad: Koninklijke Vermande, pp. 71-80, 1994 ISBN 90 5458 031 3.

More information about the JURIX foundation and its activities can be obtained by contacting the JURIX secretariat:

Mr. C.N.J. de Vey Mestdagh
University of Groningen, Faculty of Law
Oude Kijk in 't Jatstraat 26
P.O. Box 716
9700 AS Groningen
Tel: +31 50 3635790/5433
Fax: +31 50 3635603
Email: sesam@rechten.rug.nl

THE FORMALIZATION OF RETRIEVAL AND ADVISORY SYSTEMS

P. MARIANI, D. TISCORNIA and F. TURCHI
Istituto per la Documentazione Giuridica - C.N.R., Firenze, Italia

Summary

The paper deals with the characteristics of the legislative language from the point of view of its formalization. A description of a conceptual retrieval system, particularly focusing on the concepts organization and on their function in the process of query construction is given. It is followed by a description of an advisory system. The analysis will be limited mainly to the distinction between the representation of technical concepts (frames) and normative processes (rules).

1. Introduction

A great deal of literature in the field of Artificial Intelligence (AI) is dedicated to the development of a methodology for representing knowledge and for theorising about the basic requisites for expressive and efficient representation languages.
Many scientists agree on the fact that general-purpose knowledge representation should allow for a sound and complete knowledge retrieval. In logical terms, this means that every answer is always entailed by the axioms (the knowledge), and, at the same time, that it should be possible to compute all entailed propositions. This requirement of strictness, at least virtually, is mainly met by an explanation and precise definition of the concepts, most frequently organised in taxonomic trees. These enable the system to understand new concepts and to carry out logical operations, such as subsumption, implication and equivalence recognition.

To avoid difficulties on the level of computational efficiency, it has been proposed to restrict the language both with regard to the number of defined linguistic entities and to the classifications. The restrictions considerably limit the expressive capacity of the language [Brachman, 1985][Doyle, 1991].

In general, the trade-off between expressive richness and computational efficiency is difficult to handle. Therefore, it seems hard to reach a general-purpose knowledge-representation methodology. A price remains to be paid in expressiveness for efficiency or vice versa. In this article, we intend to contribute to the theoretical debate by describing concrete applications and by focusing on some specific aspects:

* the particularities of sectorial languages (in our case, legal language) compared to a global approach;
* the incidence of the kinds of utilization in the definition process of a representation language;
* the differences resulting from the use of different tools: in the first project the semantic network is formalized in first order logic and implemented in Prolog; in the second, we have chosen a pre-constructed shell with, the possibility of using different formalisms within the same structure.

The two applications described below are in the field of the Knowledge-Based Systems, we use the terminology that has become traditional [Sergot, 1988], within the sphere of systems that use formalized knowledge for aiding in decision-making, planning, problem solving, diagnosis, etc.. The projects have the same legal domain of application: environmental law; in fact, both are based on knowledge taken from the European Community and Italian legislation on the environment. However, the systems have different objectives: in one case the objective is decision support, in the other it is the retrieval of documentary information. The uniformity of matter allows to illustrate the differences and similarities between the two projects and, therefore, to make general considerations about the knowledge-representation methodology, especially about the role carried out by the representation of relevant concepts (legal or non). We will focus on the following aspects:

* the need to integrate the syntactic/structural level with a conceptual approach at the semantical level;
* the correlated need to organize these concepts in frameworks that permit the program to interpret them;
* the need to deal with the accessibility of the data and to ensure wide navigation within it;
* the distinction between kinds of knowledge: terminological and relational knowledge; technical and legal knowledge; dynamic and static knowledge;
* the need, even in the knowledge representation phase, to take in consideration into the uses of the system, for example: the amount of detail will be greater in an advisory system than in a retrieval system.

2. The formalization of legislation

On the legal level, the structuralization of knowledge is based on the definition of criteria for extracting the relevant concepts from legislative texts and for organizing them. It means associating the objects of knowledge with the terms that express them and correlating the concepts amongst themselves. For the same concept to be placed in relation to other concepts, or other objects of the domain in reference to each of its relevant characteristics, one of the main problems seems to be that of being able to explicit all the properties necessary for generating more complex knowledge. Methodologically, the most difficult problem is met when identifying *primitive* concepts. Given a knowledge domain, all elements (which could be defined as *atomic*, from a semantic point of view) are assured to be able to generate the knowledge required for understanding the domain. A second difficulty arises in identifying the level of detail that we plan to go into. A third difficulty is linked to the formal definition (above all on a semantic level) of the relations.
In identifying the concepts, or classes of concepts (types) relevant for a legal domain, four pragmatic criteria could be of assistance:

* to identify the objectives of the legislator i.e., the object of the norms);
* to take into account the user's identities (lawyers or laymen) and the purposes of utilization;
* to expand the semantic links between the legal terms used by the legislator; the links between the legal concepts can be found not only in the legislative text but also in other written normative sources which may originate from the systematic development of legal authority;
* to base the definition of the entities, their attributes and the relations between them on legal theory and case-law.

In fact, the law has the advantage of offering a consolidated (even if not unambiguous) theoretical framework of a systematic structure; four large sub-classes can be distinguished: subjects, objects, acts, facts. Objects and facts represent the more real-world aspects, while subjects and acts represent the more strictly human aspects. Subject and object are spatial phenomena: they constitute the point where subsequent fact situations are linked (for example: the transferring of property); while facts and acts belong to the category of temporal phenomena: they characterize the facts of a case that are connected and differentiated in time.

In the logical structure of the norm this distinction of legal phenomena is translated into an identification of subjects and objects which are mapped onto individuals and, therefore, onto logical subjects. Acts and facts are relations or properties attributed to them and, therefore, mapped onto logical predicates. We understand the concept of individuals as denoting classes of elements within the logical universe [Falzea, 1967]. The above distinction resembles the specification that AI theorists assume to be present between terminological knowledge (entities, concepts) and relational knowledge (meaning or factual relationships between entities); in passing, we refer to a more technical distinction between *concepts* and *roles* [Brachman, 1984]. The knowledge organization of the system described in section 2 uses this distinction.

3. The retrieval system

3.1. The domain: description of the subject matter

In Italian law, the environmental domain and, in particular, the environmental law consists of a considerable number of provisions. They are often heterogeneous and mostly presented in different levels, so that they cannot always be readily identified in relation to the domain under investigation. Sectorial and fragmentary legislation directed at regulating only some of the more frequently polluting factors is indeed often involved. Furthermore, the most recent trend in legislation making is to integrate and overlap prior legislation that is still in full force. It often contains regulations that are fundamental as well as sectorial about phenomena regarding the environment.
This situation of the abnormal mixture of laws has been aggravated by an excessive proliferation of central and peripheral institutions with competence at different levels (State, Regions, Provinces, Municipalities, U.S.L. etc.). For the environmental protection this results in uncertainty about who has to exercise control and has to take action. Poor legislative drafting makes a piece of legislation not seldom incomprehensible. As a solution one might think of widening the domain, resorting to case law and legal authority in order to be able to place the rules of individual sectors within a proper perspective.
In the phase of designing the prototype, we restricted ourselves to a subset of the domain: noise pollution.
From the point of view of document retrieval systems, the situation cannot be said to be satisfactory. There are many data banks with an enormous number of documents on-line such as the data banks in the Italgiure system[1]. Their environmental law information is spread over numerous different files which are located at different documentation centres. The problem is not so much the lack of information but the access to the documentation stored. The user who queries legal data banks, therefore, faces the following problems:

* multiplicity of data banks;
* structural diversity of the stored data;
* different link-up procedures;
* diversity of data-bank search languages;
* diversity of tools for conceptual retrieval (thesauri, classifications).

The aim of the retrieval project is to build a system that aids user-computer interaction during the retrieval process of normative sources in environmental law. It has, therefore, to make it possible for the user to search the documentary units in different data banks in a logical homogeneous form, supporting the user in the query formulation and in the definition of the search strategy.

3.2. Functionality and structure of the system

In order to define the required functions of the system, it is useful to analize the task carried out by a human intermediary who helps the user to satisfy his information needs. In brief, he carries out the following tasks:

a. he understands the user's request: he talks to the user until any ambiguity is eliminated;
b. he formulates the query and translates it into the formal database query language and defines a search strategy;
c. he analyses the result: he shows the user the retrieved information and points out which parts of the query are still unanswered. If need be, he checks why certain answers are wrong.

The proposed system, called CABALA (Consultazione Assistita di Basi di Dati di Leggi Ambientali) [Ciaramella, 1989] [Guidotti, 1990a] [Guidotti, 1990b], has four main feautures, which enable it to perform a few parts of the above listed tasks:

* navigation on a semantic network;
* handling of the dialogue with the user;
* query construction;
* definition of a search strategy.

The main aim is to display the conceptual structure of the data described by a semantic network, and to assist and guide the user in "navigating" it so that he can formulate his query properly. The query is then translated into a suitable form for the specific information retrieval system involved. CABALA can be divided into three main logically independent parts, which communicate among themselves only by exchanging messages:

* Query Generator;
* Data base Query Manager;
* Data Bases.

For the purposes of this paper, we mainly focus on the Query Generator: its objective is to enable the user, regardless of his experience in the domain, to formulate a valid query from both the legal and common sense point of view. Once the user has formulated his query, this is translated into an intermediary language and both its correctness and efficacy are evaluated. The Query Generator is, therefore, subdivided into three modules (see fig. 1):

* the Navigator
* the Query Constructor
* the Query Evaluator

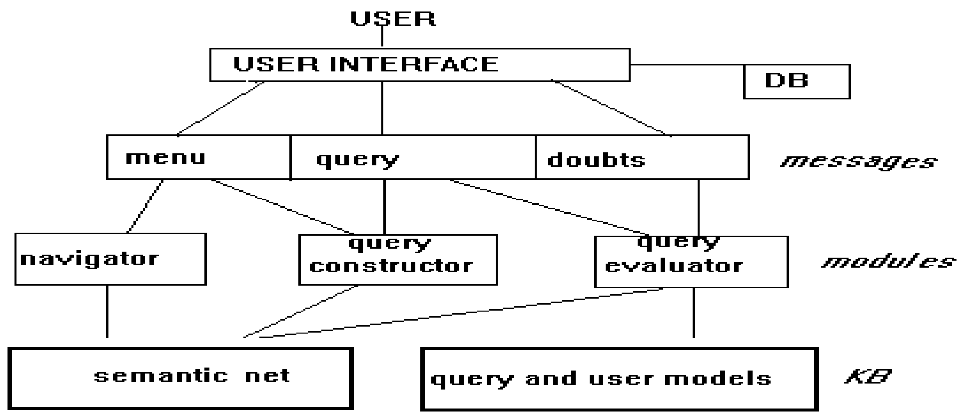

Figure 1: The three modules of the Query Generator.

3.3. Knowledge formalization

Both the Navigator and the Query Constructor utilize knowledge formalized in a semantic network : it has been created by extracting a series of significant terms from legal texts, representing nodes in the network, and defining two kinds of relations between nodes.
The Broader Term (BT) relation is transitive; it lays down hierarchies (Trees) among terms, each belonging to a specific group that arises in the description of the hypothetical case or that defines legal concepts pertaining to the domain. In the field of noise pollution the following BT hierarchies were identified: Sources of emission, Types of noise, Ambits for safeguarding, Prevention and control, Bodies, Normative sources, Sanctions, and Pertinent facts.
The BT hierarchies point out the basic concepts that make up the query: the nodes of the network are linked by oriented arcs, identifying the relations existing between them. That enables the system to capture the semantic of the user's queries. Let us suppose that the user selects terms in more than one hierarchy: the path on the network that links the hierarchies univocally identify the query. For example, on the basis of the terms used and on the relations between hierarchies, it is easy to capture the meaning of the following question:

* Which judicial bodies are competent to control and prevent noise pollution?
* Has the judge the power to verify the tolerance level of noise?
* To which bodies should workers address their claims to be protected against noise pollution in the working places?

We use the generic name of Related Term (RT) to identify this second kind of knowledge: the specific meaning of the relations as far as their features concern (for example directionality, transitivity, symmetry) is based on a partial interpretation of the norms, and the case law that constitutes the databases. Two examples of this kind of knowledge are:

a. *In the case of noise pollution, only fines and arrest can be invicted as sanctions by the judge*.
   Moreover, this relation enables the representation of general knowledge of law and legal authority of the type;
b. *the ordinary courts inflict criminal sanctions; the administrative courts inflict administrative sanctions*.

3.4. Use of the semantic network

The Navigator allows the conceptual structure of the documents stored in the databases to be analized: the user specifies his query from the menu in which he selects some pre-defined terms (items). The menus are derived from the hierarchies defined by the BT relation, which means that specifying the query is the same as selecting nodes in those hierarchies.
In this way, the set of hierarchical nodes performs the role of the system's vocabulary. This considerably simplifies the user-system interaction: the system never has to analyze an unknown term as may occur by freely interacting in natural language, but the user is forced to express his query in the terms of the system's vocabulary. This limitation is counterbalanced by the fact that, using hierarchically organized terms, the user is offered a classification of the domain which acts as a map of the conceptual structure of the database. We leave the user the possibility to select the hierarchies without any predefined order informing him, however, about the hierarchies correlated to the selected one.
The RT relation is utilized by the Navigator for displaying the related terms and for guiding the user in formulating his query. It should, also, be noted that additional information is given by the absence of a RT relation between two selected terms. In fact, if two terms belong to hierarchies having RT relations but they are not linked by a RT relation or do not belong to sub-trees whose roots are linked by RT arcs, an error in the

selection can be assumed and the user can, therefore, be informed about it. This is the case, for example, where the user selects "mayor" and "criminal sanctions".
In the Evaluator implementation, we thought that it was more suitable to use simple heuristic criteria like those generally followed in developing the system, which can be set up in an easier way than statistical models. The latter are more reliable but also more complex and need a large amount of data.
The Evaluator has the task of verifying the validity and the efficiency of the query. Therefore, the first thing that the Evaluator has to do is to verify whether it is possible to find a path linking the selected hierarchies; if not, the Evaluator suggests how to complete the path. Anyhow, it is always possible to force the system to process the query.
The Query Constructor accepts the selection about the hierarchies made by the user and transforms it into a query for the textual data base.

4. The advisory system

4.1. Object, aims, general features

The objective of the SEDAM Project[2] (*Sistema Esperto in Diritto dell'AMbiente* - "Expert System in Environmental Law") is to build an expert system that, making inferences from the knowledge of legislative sources on air pollution, is able to solve problems in the area, adding the pertinent documentation to the proposed solution (sections of statutes 'applied' for reaching a solution, notes on case law and legal authority). We chose this domain both because of its social importance and because of its legislative complexity, above all in the application phase.
The system contains knowledge including all national legislation on air pollution; amongst the most recent legislation. All the pertinent European Community legislation and provisions on motor vehicles and electric power plants are also included along with the legislation of the Tuscan Region on the matter.
On a legislative level, the material has been organized around the basic concept of *emission limits*, that is the maximum values permitted for the emission of polluting substances. These values are obviously established on the basis of the type of substance, type of plant, power of the installations, etc. The other cardinal concepts are that of *authorization* in exercising an activity, and that of *penalties* that can be imposed where limits have been infringed. The basic problems the expert system has to solve involve parameters relative to lawful emissions. Furthermore, the expert system informs the user about the requisites for setting up or modifying industrial plants. It also illustrates the procedures to be followed in obtaining authorizations, the methods used for sampling and checking emissions, the kinds of penalties, the competent authorities, and the legislative sources.

4.2. The formalization criteria

The main differences between the retrieval and the advisory system lie in the facts that (1) the amount of information (supplied by the user) about the *fact situation* is obviously greater in the advisory system than in the retrieval system and, (2) likewise, the knowledge *a priori* embedded in the system must be complete and articulate in the advisory system so that the reply can be inferred.
All technical aspects of the domain cannot be ignored but, indeed, they become decisive both from the point of view of enabling the situation under examination to be described and from the point of view of giving the specific values provided by law for every fact situation. Furthermore, the differentiation of plants into "new", "existing" and "modified" is fundamental on a legislative level - an aspect that is, therefore, transversal to the description of the type of plant.
Also, with regard to types of chemical substances that constitute the emissions, the law classifies them on the basis of their physical state (dust, gas, vapour) and of their danger.

The emission threshold limits permitted by law are, therefore, based on different parameters: types of plants, state of the plants, types of substances. Values are set for every substance on the basis of a pre-established reference value.
On the representation level, the importance of technical knowledge compared with more strictly prescriptive knowledge should, therefore, be taken into account. The lawfulness of the described situations, in practice, is solely determined by the fact that the actual emission values and the normative limits agree. We have, therefore, opted for a detailed representation of factual aspects, focusing our attention on the phase of information acquisition from the user and on adopting inheritance techniques for reducing and structuring the knowledge.
Therefore, considering the relevance of technical knowledge, the problems are more related to questions of computational efficiency than to semantic issues. Instead of the semantic network formalism, that allows description of conceptual entities (and the further distinction between terminological and relational aspects), we have preferred to use the formalism of frames and to focus on the factual characteristics of the knowledge.

4.3. Flex

The shell we have used for developing the expert system is Flex, a toolkit for expert systems[3] that offers different integrated formalisms and inference engines, based on the Prolog language. The codification language is similar to the common English language, in which frames, instances, attributes and production rules can be expressed, and that allow the definition of personalised Prolog predicates. The expert system structure reproduces the classical architecture of these programs: knowledge base; inference engine; user interface.
Interaction with the user occurs through a system of windows with pop-up and pull-down menus.

4.3.1. Use of frames

Frames have been used for representing technical knowledge: classes of chemical substances, types of industrial plants, categories of production processes. The framework has permitted us to express complex concepts and types of different information linked to every entity: for example, for every type of emission, there are limits that vary according to each chemical component, rate of concentration, classification on the basis of their danger and physical state, etc.
The frames' structure is particularly suitable for describing classes of objects and abstract concepts. By employing database terminology a *frame* may be seen as a *record* and the *slots* as *fields of a record*; every slot has three components: an *attribute* or, in other words, the name of a characteristic applicable to the slot, the *default*, that is, the omitted value for the attribute and the *current value* of the attribute. The advantages in using frames for knowledge representation lie mainly in their flexibility in processing complex information and in their capacity to incorporate procedural aspects: a set of conditions (constraints) that are to be verified in the case of updating or in the case of access to the slots can be associated with every frame. Moreover, frames are particularly suitable for developing forms of reasoning by default, utilising, in our case, for inheritance, for example, in classifying plants, the unchanged emission limits in the subspecifications of the type of plant.
As an example of a frame in Flex, the structure below expresses a segment of the knowledge about emission, in three levels, to the lower of which are associated the normative limits (the keys word are in bold characters):

frame emissions
frame carcinogenic_substance **is a kind of** emission
frame group_1_carcinogenic **is a kind of** carcinogenic_substance
default threshold_value **is** "for 0,5 g/h: 0,1 mg/m3 maximum flow ".

The user selects substances trough the following structure:

group group_1_carcinogenic
"asbestos"
"benzoine"
"chromium"
"cobalt"

The system needs a further information for retrieving the correct value:

instance a_s **is an instance of** group_1_carcinogenic;
name **is** 'asbestos'.

Through the distributed information property, typical of frames, it is possible to obtain the values relating to each selected substance: when the norms do not establish a specific value, the system will inherit it from the hierarchy.
The system shows to the user the threshold value, and compares it with the actual values supplied by the user.
In the same way, the system is able to find the values of the specific substances relating to the type of plant: for every ancestor of the selected plant (the plants are, naturally, defined with frames) the default value is determined and it is compared with a reference value by crossing, in a natural way, all the branches of the tree of the frames from one specific node down to the root.

4.3.2. Use of rules

The strictly prescriptive part of the knowledge (granting of authorizations, checking of the correspondence between values, inflicting of penalties) has been formalized using production rules. This means hypothetical sentences of the following kind:

if<condition>*then*<consequence>

Such a structure has the following characteristics:

* modularity: every rule represents a small part of the relatively independent knowledge;
* ability to increment: new rules (new knowledge) are easily added;
* transparency: ease in explaining how the system works, or the rules facilitate the user's replies to questions like "How was a certain conclusion reached", or "Why was this information asked for".

Two types of rules can be defined in the Flex language: one is based on Prolog clauses and uses the *backward-chaining* method of reasoning that is typical of Prolog; the other is based on a typical production rules format, based on the *forward-chaining* method of reasoning.
A feature of a rule system used in the forward-chaining mode is to permit more than one proposition in the *then* part of the conclusion. The *ruleset* structure that activates the forward-chaining method consists of:

a. determining the rules that can be used in selection (agenda);
b. defining a method for choosing the rules: Flex provides three pre-defined methods (first- come-first served, conflict-resolution and conflict-resolution with threshold value) and one that can be defined by the programmer;
c. establishing the method for updating the agenda.

Furthermore, an explanation of the reason why it was applied can be associated to every rule, through a link up with a system for consulting the files.

5. Conclusion

A well-known characteristic of the environmental domain is the coexistence of technical data (types of substance, types of plants, sources of emissions, etc.) and legal data (competent institutions, types of penalties, legislative sources, etc.). As technical knowledge may be considered static, at least within a limited legislative sphere, it was used for its representation formalisms, like semantic networks and frames, that allow taxonomic relationships and inheritance between the concepts to be expressed easily. Furthermore, these enable knowledge for branches of a domain to be modularised, an aspect that should not be ignored if we consider that the more strictly legal part of environmental law is almost common to all the fields of application: it may, therefore, prove advantageous to keep knowledge on the domain separate from it, which would allow the knowledge base to be added to without retouching the pre-existing base.
Therefore, the traditional distinction between *terminological knowledge* and *assertional* (*or relational*) *knowledge* is here taken into consideration on a level of a rough, not "granular" analysis; this means that it is possible to make it coincide with the distinction between *static* or *"stable" concepts* (for example: legislative definitions, judicial bodies, types of plants, classes of substances), and *contingent legal situations*, linked to factual data and evolving dynamically (process for the granting of an authorization, infringement of regulations, etc.). Such a distinction, facilitated in the second project by using the different structures of the tool, would appear not to be necessary in the first project, where the representation of factual and contingent aspects does not occur: in fact only terminological knowledge is sufficient for processing information of a purely linguistic nature, where the semantic aspects (the arcs that link the nodes of the network) come within the aim of completing and customising the conceptual precision of the documentary retrieval.
As a consequence, in the retrieval system it is satisfactory enough to maintain the formalization of the conceptual aspects at a surface level; that is, the system is not based on a "deep model" of knowledge, as it is traditionally required in the field of "conceptual retrieval systems" [Hafner, 1991], where the concepts understanding make it necessary to complete the knowledge on the stored data with knowledge about the world.
Also completeness in defining the concepts does not seem to be a necessary requisite in both cases: the objective is not to accept and understand new concepts (by comparing and placing them in complete taxonomic trees) but, rather, to help the user make correct choices in expressing his query through formalized concepts.

6. Notes

1 The Italgiure is the main italian legal information system.
2 The project was developed at IDG by a research team made up of A. Cammelli, P. Mariani, F. Socci, D.Tiscornia, F.Turchi.
3 Distributed by LPA Ltd. London.

7. References

This article may only be cited as: Mariani, P., D. Tiscornia and F. Turchi, Contribution to a Methodology for the Legal Knowledge Representation: Two Project in Environmental Law Domain. In: Grütters, C.A.F.M., 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 & Law, JURIX'92*, Koninklijke Vermande, Lelystad, NL, 1992.

[Brachman et al., 1983] Brachman, R.J., R.E. Fikes and H.J. Levesque, Krypton: a functional approach to knowledge representation. In: *IEEE Comput.*, 1983, no. 16, pp. 67 -73.

[Brachman & Levesque, 1985] Brachman, R.J. and H.J. Levesque, A fundamental Tradeoff in Knowledge Representation and Reasoning. In: Brachman, R.J. and H.J. Levesque (eds), *Reading in Knowledge Representation*, Kaufmann, Palo Alto, CA, 1985.

[Ciaramella et al., 1989] Ciaramella, N., P. Guidotti, L. Lucchesi, P. Mariani, M. Ragona and D. Tiscornia, Un sistema di reperimento concettuale per archivi sul diritto dell'ambiente. In: Martino, A.A. (ed.), *Sistemi esperti nel diritto*, Cedam, 1989, pp. 491-509.

[Doyle & Patil, 1991] Doyle, J. and R. Patil, Two theses of knowledge representation: language restrictions, taxonomic classification, and the utility of representation services. In: *Artificial Intelligence*, 1991, no. 48, pp. 261-297.

[Falzea, 1967] A. Falzea, Voce "Fatto giuridico". In: *Enciclopedia del Diritto*, Giuffr, vol. XVI, 1967, p. 942.

[Guidotti et al., 1990a] Guidotti, P., L. Lucchesi, P. Mariani, M. Ragona and D. Tiscornia, An Intelligent Interface to Data Bases on Environmental Law. In: *Proceedings of the Congreso Internacional de Informatica y derecho*, Buenos Aires, 1990.

[Guidotti et al., 1990b] Guidotti, P., L. Lucchesi, P. Mariani, M. Ragona and D. Tiscornia, A simple intelligent Interface to Data Bases on Environmental Law. In: *Database and Expert Systems Applications, DEXA '90*, Springer Verlag, 1990.

[Hafner, 1991] Hafner, C.D., Artficial Intelligence and Law: Opportunities and Challenges, Tutorial presented at *Third International Conference on AI and Law, ICAIL'91*, Oxford, 1991.

[Sergot, 1988] Sergot, M., Representing Legislation as Logic Programs. In: *Machine Intelligence*, Oxford University Press, 1988, no. 11.