Use of case-based reasoning in the domain of building regulations

Soon-Ae Yang† **Dave Robertson**†
(soonae@aisb.ed.ac.uk) (dr@aisb.ed.ac.uk)
†Dept of Artificial Intelligence
University of Edinburgh
80 South Bridge
Edinburgh EH1 1HN, UK

John Lee‡
(john@caad.ed.ac.uk)
‡EdCAAD, Dept of Architecture
University of Edinburgh
20 Chambers Street
Edinburgh EH1 1JZ, UK

In Proceedings of the Second European Workshop on Case-Based Reasoning, 7-10 November 1994, Chantilly, France.

Use of Case-Based Reasoning in the Domain of Building Regulations

Soon-Ae Yang[1], Dave Robertson[1], John Lee[2]

[1] Dept of Artificial Intelligence, University of Edinburgh,
80 South Bridge, Edinburgh EH1 1HN, UK
[2] EdCAAD and Human Communication Research Centre,
Dept of Architecture, University of Edinburgh,
20 Chambers Street, Edinburgh EH1 1JZ, UK

Abstract. In traditional legal decision support systems, it has been regarded as natural to represent statutes in terms of decision rules and to link these to a separate case-based reasoning system for handling precedent. Statutory legal rules used in these systems are formal and prescriptive. Building regulations in Scotland are part of statute law and constitute part of a legal system together with case histories. In recent years, the regulations have been becoming less prescriptive and more emphasis has been put onto the interpretive use of the regulations. In developing a system which can assist domain experts in interpreting the regulations, this trend has presented us with difficulties in employing this traditional approach and has led us to a unified case-based model of the regulations and case histories. In this paper, we first describe the characteristics of the regulations and the activities involved in this domain. Second, we explain the reason why we abandoned the traditional approach. Third, we describe the system which has been developed using this case-based model.

1 Introduction

Legal and regulatory reasoning is one of the AI application domains where substantial amounts of research have been done over the last twenty years [7]. The law contains "legal rules", which state what should be done in certain circumstances in the context of a legal system. Most of these are expressed in the form of statutes and cases. Building regulations in Scotland are part of statute law and constitute part of a legal system together with case histories.

Since it is normal to view statutory regulations as "legal rules", it is also natural to expect these to be most easily expressed formally in terms of decision rules (perhaps in some propositional language). However, it is becoming increasingly common to make the regulations less prescriptive, with the text of the regulations merely providing a general framework within which the interpretation of the regulations is established through precedent. In these types of system the legal rules do not in themselves allow any decision to be made. Much greater importance is attached to the means by which precedent is accessed and interpreted. A good example of this approach is found in the building regulations

of the Scottish Office. These will be used throughout this paper as a concrete example of the more general application.

2 The Domain of Building Regulations

The Building Directorate of The Scottish Office is responsible for the drafting and maintenance of building regulations—the Building Standards (Scotland) Regulations 1981 [9]—which set out the statutory requirements affecting building design and construction in Scotland. The Building Directorate is also responsible for the administration of the appeals procedures under the building control system in Scotland.

Building regulations have the generic objectives of safeguarding public health and safety and of conserving energy in the built environment. These objectives express government policy and the principles which bring building regulations into existence. These issues or principles are translated into more detailed issues in the sixteen subject areas: adequate ventilation, drainage and sanitary convenience to ensure the health of people, fitness of materials for durability to ensure the safety of people, structural stability in case of fire to ensure the safety of people, adequate means of escape from fire to ensure the safety of people in case of fire, safe means of passage, etc. These issues or requirements are very general and expressions are highly abstract. The bounds of applicability of this requirement are, intentionally, left unclear at this point. These requirements can be met by compliance with the relevant regulations and provisions set out in the Building Standards (Scotland) Regulations. Requirements stated in these relevant regulations and provisions are a refinement of more general and abstract requirements described above and can be regarded as the description of how those requirements should be interpreted. For example, Regulation D7 (see Fig. 1) requires that building components be constructed of non-combustible materials to maintain structural stability in case of fire.

Requirements are specified only for certain aspects or features of buildings that are related to a particular generic objectives. Other aspects or features unrelated to this particular generic objective are disregarded. In the example in Fig. 1, requirements to achieve structural stability specify only non-combustibility of materials used for particular building components such as compartment floor, floor of any landing or passage within a stairway enclosure and stair within a stairway enclosure. Other factors, for example the thermal transmittance value of materials and the shape and position of stairs are regarded as irrelevant in the subject of structural stability. These requirements do not always apply to every possible building category and different requirements are sometimes specified for different building categories. For example, as shown in Fig. 1, "the floor of any landing or passage within a stairway enclosure" in "a house in occupancy subgroup A2, not being a flat, not more than three storeys in height" does not need to be constructed of non-combustible materials. Cases are also described in the same manner. They do not provide a design of a whole building but provide only the description of aspects or features of a building related to a particular

D7 Requirements as to non-combustibility
Every part of a building specified in column (1) of the following table shall be constructed of non-combustible materials, subject to the exceptions if any specified in relation thereto in column (2): ···

Table to Regulation D7	
Requirements as to non-combustibility	
Parts of a building required to be constructed of non-combustible materials	Exceptions
(1)	(2)
Compartment floor	
Floor of any landing or passage within a stairway enclosure provided so as to comply with Regulation E10	The floor of any landing or passage within the stairway enclosure of a stair in a house in occupancy subgroup A2, not being a flat, not more than three storeys in height
⋮	⋮
stair within a stairway enclosure provided so as to comply with regulation E10 or unenclosed external stairway referred to in regulation E10(1)	1. Any handrail on such a stair 2. Any stair in a house in ···
⋮	⋮

Fig. 1. Regulation D7

subject area in question.

The relevancy between issues and requirements, which plays a central role in interpretation of the document, is sometimes clearly expressed in the document but in some cases may rely on experts' experience in applying requirements to particular problems.

The interpretation of written requirements also changes through time. At the time of legislation, it is impossible to anticipate all possible circumstances in which a case may arise after legislation. New legal rules are set out in unexpected circumstances or slightly different circumstances and existing legal rules may be relaxed in exceptional circumstances. If it is unreasonable for specific requirements to be applied and these requirements can be met in a less onerous way, existing requirements may be relaxed. In this sense, cases may set out new legal rules but legal rules set out in case histories may or may not apply to new cases. As legal rules from case histories accumulate, they are reviewed when revising statutory regulations. The rulings made in case histories may be generalised and included in the new version of regulations.

An application for relaxation of particular requirements is first submitted to the Local Building Control Authority and if there is disagreement about interpre-

tation between the applicant and the Authority the appeals procedure follows. The argument surrounding each appeal decision is often complex, involving the balancing of a variety of assertions by the parties concerned. Experts make a decision after (if necessary) consulting all information relevant to the given case. In addition to statute and case histories relevant to the given case, experts should refer to all information available to provide justification for their decision within the legal system in which they work. This includes government policies, principles taken for granted in the given field of law, and other related law. It would be ideal to build a knowledge base which includes all this information but that would be too ambitious. The problem is therefore to support the consultation of these loosely expressed legal requirements without being forced to make explicit all the information needed to arrive at a decision.

We do not employ a rule-based approach in our system and employ a purely case-based approach for both statute and precedent. In the next section we explain the reasons why we avoided a rule-based approach.

3 Why Not a Traditional Legal Decision Support System in This Domain?

Legal rules are the basic unit of knowledge to be represented in the knowledge base. Legal rules are requirements, which describe a required design *solution* in certain *circumstances*. Given this description of legal rules, conventional if-then rules seem a natural choice for representing legal rules—for example, in the form:

if *circumstances* then *solution*

For example, Regulation D7 can be described in this form as shown below. For ease of understanding, *circumstances* and *requirements* are given not in computer-executable form but in approximate English.

if *the type of building is any and there is a stair in an enclosed stairway used as an exit*
then *the material used for the stair should be non-combustible*

Then some form of rule interpreter might be used to derive information from legal rules represented in the knowledge base. This approach has been taken in Waterman and Paterson's LDS [17], Sergot et al's representation of the British Nationality Act 1981 [10, 11] and Susskind's system [13]. The Scottish Office's Building Directorate has also experimented with this approach in their in-house prototype systems [12]. Legal rules used in these systems are formal and prescriptive. However, the regulations are becoming less prescriptive and more emphasis is put onto the interpretive use of the regulations. This presents us with several difficulties in employing a traditional decision support system approach.

1. Legal rules are described at different levels of detail.

In addition to legal rules which are prescriptive and state requirements at a very detailed level, such as Regulation D7, there are legal rules which are less prescriptive and which explain why these detailed requirements should be satisfied in the context of the legal system.

Legal rules are also independently described at various levels of detail. Legal rules at the upper level describe requirements using more abstract terms which have broader meaning. Terms like "safety in case of fire" and "adequate means" are used at this level. Then one or more legal rules are introduced to provide various ways of interpreting these abstract legal rules. Refined legal rules are described using more detailed and concrete terms. For example, what "safety in case of fire" means in the context of achieving "safety" is defined in terms of "structural stability", and then "storeys", "compartment walls", "compartment floors", and so forth. In other words, legal rules at lower levels are a refinement or specialisation or decomposition of legal rules at upper levels.

2. There exists a terminological gap between legal rules and cases.

Although appropriate rule structuring might overcome the problem of different levels of detail in statute, the language used in statutory regulations is generally at a more general level of description than that used in describing the objects (buildings) to which the rules apply. The description of the circumstances and required solutions in statutory regulations is the description of requirements generalised from many particular circumstances and solutions in these particular circumstances. Cases arise in one of those particular circumstances and propose one particular solution. For example, in cases, the type of building may be given as "flat" or "house", which are of type "dwelling", or instead of "rooflight" the term "window" may be used. It would be possible to bridge these terminological gaps by introducing domain-specific "bridging" rules, but it is not always possible reliably to bridge this terminological gap in this way, as we demonstrate below.

3. The meaning of terms is often in debate.

It is sometimes a matter of debate as to what the regulations mean by basic concepts. For example, does "floor area" include or exclude raised platforms? or is embankment "a place of safety"? This is one of the most common reasons for appealing against the regulations and seeking relaxation, and there are prolonged arguments as to the meaning of terms in the context of the given circumstances in many appeal cases provided by the Scottish Office's Building Directorate. Although it would be possible to encode in rules how these terms are interpreted in particular circumstances, there would remain the serious practical problem of maintaining consistency across the vast collection of these "bridging" rules, given that interpretation of the regulations constantly changes over time and there is always the possibility of conflicts between these bridging rules.

4. New legal rules are constantly emerging.

Even if the above theoretical and practical objectives could be addressed, we would still be left with the problem of accommodating new, unexpected or exceptional circumstances within the decision support system, for example if someone invented a new type of building material.

5. Legal rules are the generalisation of the rulings made in case histories.

One possible way of solving these problems (1-4) is to use the rule-based model for statutory regulations and to use the case-based model for case histories such as in CABARET [8], GREBE [1] and PROLEXS [16]. These systems assume that legal rules in statutes are open-textured and may (to a limited extent) conflict with each other but they are formal and fairly stable. The problems with open-texturedness and conflicting rules are compensated for by referring to cases.

But we take a different view of statutes and cases. Legal rules in the statutes are the result of accumulation and generalisation of the rulings made in case histories. A case must comply with legal rules in the statutes which are applicable. In this sense, legal rules in the statutes act as "strong" constraints on new cases. Cases provide a way of interpreting a particular set of legal rules in the statutes under particular circumstances, in other words more detailed legal rules than those in the statutes. There may be similar cases and it is possible to derive more generalised legal rules from those similar cases. The interpretation of legal rules in the statutes changes over time, legal rules set out in case histories may or may not apply to new cases. In this sense, legal rules set out in case histories act as "weak" constraints to new cases.

Once we view legal rules in this way (i.e. legal rules are generalised cases) legal rules from both statutes and cases can be represented using only a case-based model. In the system described in the next section, cases are kept in a separate case library, and legal rules, i.e. the generalised cases, are represented as *models*, "strong" models when a model represents legal rules from the statutory regulations and "weak" models when a model represents legal rules set out in a case. Interpretation hierarchies which exist among legal rules can be represented as abstraction hierarchies of models. Each case is relevant to a model or models and links between cases and models can be established using the notion of relevance. There are various matching algorithms which can be used to access models at various levels in the abstraction hierarchies. Models in the abstraction hierarchies are loosely connected to each other so that new legal rules can be added as "weak" models to the existing abstraction hierarchies even when new "weak" models are somehow inconsistent with the existing models.

A prototype system which can assist domain experts in the appeal process has been developed using this approach. The system does not make any decision about relaxation but is intended to provide the experts with relevant information, i.e. relevant legal rules and cases, so that they can arrive at an informed decision. The next section describes each component of the system briefly. [3]

[3] A more detailed description of the system can be found in [18].

4 The Overview of the System

There are three knowledge bases: the Model Knowledge Base, the Case Library and the Domain Knowledge Base. Legal rules are represented in the Model Knowledge Base. In addition to legal rules, background domain knowledge must be used to bridge terminological gaps between models at different levels of abstraction hierarchies and between models and cases in the retrieval and maintenance of legal rules. Cases and argumentation surrounding each case must be provided to the user. Argumentation surrounding each case is represented in the form of an argumentation network. Background domain knowledge is represented in the Domain Knowledge Base, and cases and argumentation surrounding each case are represented in the Case Library.

A case will be given as $C = (C_{id}, C_g, C_c, C_{cd}, C_{sd})$ where C_{id} is the id number of the case, C_g is the goal given in the case, C_c is the category of the building, C_{cd} is the circumstances description, and C_{sd} is the proposed solution. Given a case C, the system first retrieves relevant models by comparing the C_g, C_c and C_{cd} of the given case with models in the MKB. Once relevant models are found, the C_{sd} of the case is evaluated against the required solution of the retrieved model. The case is said to be *compliant* if the C_{sd} satisfies the required solution of the retrieved strong models.

If the case is not compliant, there is a possibility of relaxation and the system will explain relevant information stored within the system: principles or issues behind the regulation with which the case is supposed to comply and previous similar cases. The system also checks whether there is conflict between the case and existing models and informs the expert. If the expert decides to grant relaxation, new rulings set out by the case will be integrated into the existing abstraction hierarchy by (if possible) combining new rulings with relevant weak models. The case itself will also be stored in the Case Library. The link between the case and the MKB is established at this point by storing the id numbers of relevant models with other information in the Case Library.

4.1 The Model Knowledge Base

Legal rules are represented as *models* and then an abstraction hierarchy is built for each generic objective. In each abstraction hierarchy, a model M is represented as a set of six parts: $(M_{id}, M_{pid}, M_c, M_{cd}, M_{sd}, M_s)$ where M_{id} is the id number of the model, M_{pid} is the id number of the parent model in the abstraction hierarchy, M_c is the category of buildings to which the model applies, M_{cd} is the description of key features of circumstances to which the model applies, M_{sd} is the description of required solutions, and M_s is the strength, i.e. "strong" or "weak". The strength will be given as "strong" if the model is taken from statutory regulations, or as "weak" otherwise. M_{cd} and M_{sd} are described in terms of objects, their attributes and relations between objects. An example of a model is shown in Fig. 2.

Abstraction hierarchies of models are also shown in Fig. 2. The generic objectives are represented in M1–M4: all buildings should be built not to harm the

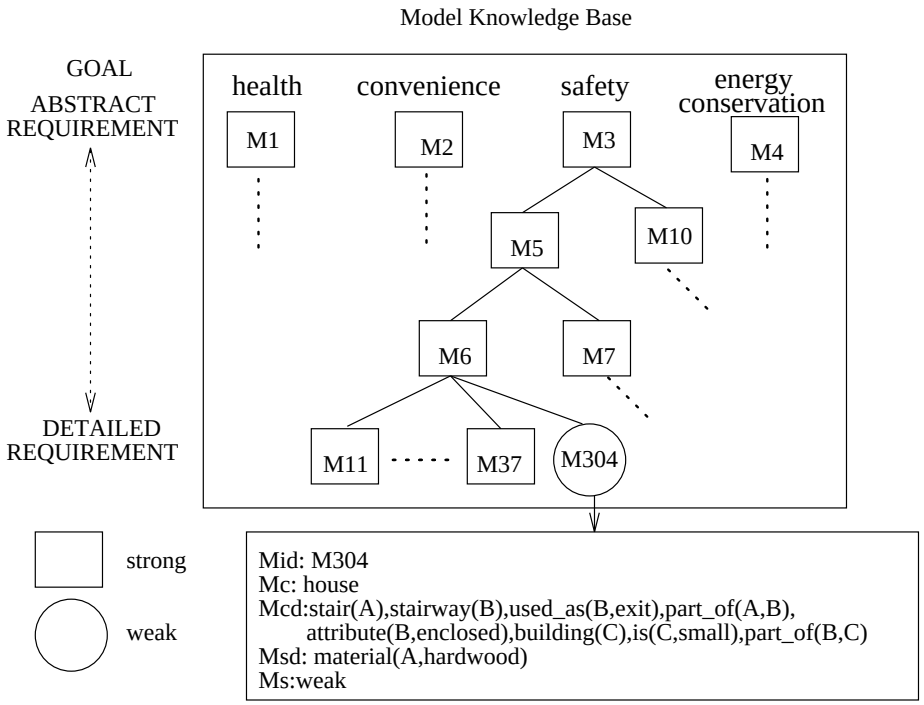

Fig. 2. Model Knowledge Base

health(M1) and to ensure convenience (M2) and the safety (M3) of people in and around buildings and the conservation of fuel and power (M4). M3 is interpreted as: all buildings should be built to ensure the safety of people in case of fire (M5) and to provide safe means of passage (M10) to achieve the safety of people. M5 again is interpreted as: structural stability should be maintained in case of fire (M6) and adequate means of escape from fire should be provided (M7). M11–M37 represent the detailed requirements to ensure structural stability in case of fire (M6) in terms of compartments, floors, stairs, walls, etc. M304 represents a ruling from a case, in which the stair of an enclosed stairway used as an exit was allowed to be constructed of hardwood. This ruling can be regarded as another interpretation of M6 and added to the MKB.

4.2 The Case Library and Argumentation Structure

The Case Library is a collection of cases processed by the system. In addition to an argumentation network, a case C is represented as a set of eight parts: $C = (C_{id}, C_g, C_c, C_{cd}, C_{sd}, C_d, C_m, C_t)$ where C_{id} is the id number of the case, C_g is the goal given, C_c is the category of the building, C_{cd} is the circumstance description, C_{sd} is the proposed solution, C_d is the decision taken by the Building

Directorate (BD), i.e. "granted" or "refused", C_m is the id number of the model which is relevant to the case, and any other information will be stored in texts in C_t. For example, Case 90 is represented as follows.

C_{id}: C90
C_g: safety
C_c: house
C_{cd}: stair(a), exit_stairway(b), is(b,enclosed), part_of(a,b)
C_{sd}: material(a,timber), fire_resistance(a,1,hour)
C_d: granted
C_m: M23
C_t: nil

Argumentation for the relaxation of specific regulations starts with usually two *positions*, one suggested by the applicant and the another one or two suggested by the Local Building Control Authority (the LBCA). Since the applicant is the player who seeks relaxation on the regulations at issue, the applicant's proposal constitutes partial compliance with the regulations. Appeal cases are submitted to the BD because the LBCA refused the relaxation and insisted on full compliance with the regulations at issue. In addition to these two *players*, there may be other players who make *arguments* to *support* or *oppose* the suggested positions such as the BD itself or the firemaster. The role of the BD is to arbitrate between these two main players—the applicant and the LBCA—and decide which position is reasonably acceptable in the context of the given case. Such argumentation is recorded in the Case Library along with cases in the form of an IBIS-like argumentation structure, which was developed from the IBIS model [6, 2] and Toulmin's model [15] of argumentation. In this argumentation structure, there are two node types, *position* and *argument*, and two link types, *support* and *oppose*. An example of an argumentation structure is shown in Fig. 3. Reasons to *support* or *oppose* a position, e.g. "full compliance", "partial compliance" and "reasonable", are enclosed in brackets ([...]) in this figure.

4.3 The Domain Knowledge Base

Knowledge in the Domain Knowledge Base (the DKB) is used in classifying a given case, i.e. to find models whose M_c and M_{cd} are most similar to C_c and C_{cd} of the given case. Specialisation hierarchies are used to represent building categories and concepts used to describe models and cases (e.g. "house" is a subclass of "dwelling", "window" is a subclass of "opening"). Aggregation hierarchies are used to represent structural knowledge about building components (e.g. "stair" is a subpart of "stairway"). Examples of specialisation/aggregation hierarchies are shown in Fig. 4.

4.4 Retrieval of Relevant Regulatory Information

Given a case $C = (C_{id}, C_g, C_c, C_{cd}, C_{sd})$, the retrieval is carried out in three steps. First, the system selects an abstraction hierarchy of the given C_g and

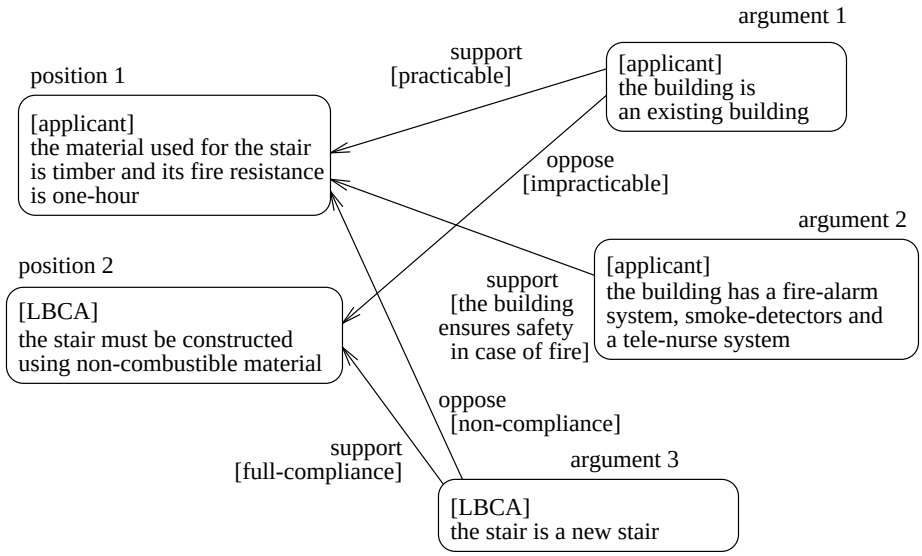

Fig. 3. Argumentation surrounding Case 90

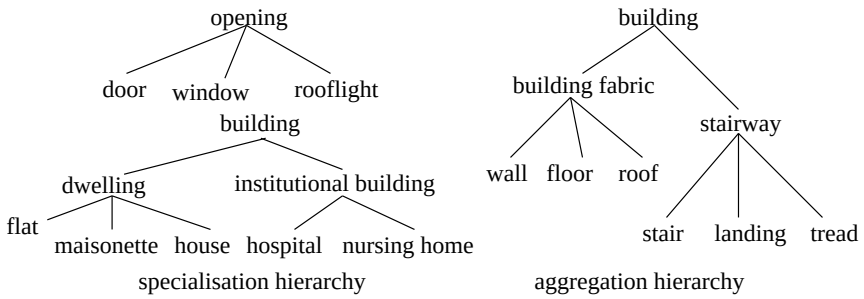

Fig. 4. Examples of specialisation/aggregation hierarchy

then selects only models whose M_c *subsumes* the C_c of the given case using the specialisation hierarchy of building categories in the DKB. M_c is said to *subsume* C_c when: a building category in M_c is the same as the building category in C_c; or a building category in M_c is at a higher level in the specialisation hierarchy than the building category in C_c.

Since models in the selected abstraction hierarchy describe legal rules using terms at different levels of detail, the second step is to find the *right level* in the abstraction hierarchy which has models comparable to the C_{cd} of the given case. This is done by selecting the level at which objects in M_{cd} are at the same level in aggregation hierarchies as objects in C_{cd}.

Once models at the right level are found, the third step is the selection of

models by assessing similarity of circumstances. The system assesses similarity by comparing the features in the M_{cd} of selected models with the C_{cd} of the given case. The system first applies the notion of structural mapping [4, 5, 14]. The system tries to find a model or (models) such that all features in M_{cd} have corresponding features in C_{cd}, in other words, for each building component in M_{cd}, there is a building component in C_{cd} which is of the same type or of the similar type with the corresponding building component in M_{cd}, and all matching building components in C_{cd} have the same or similar attributes, properties and relations with the corresponding building components in M_{cd}. Similarity between two concepts (e.g. types of building components, attributes, etc) is measured in terms of semantic distance [3], i.e. the number of links between two types in specialisation hierarchies. In the current implementation, two concepts whose semantic distance is not greater than 2 are regarded as similar. For example, the number of links between "opening" and "window" in the specialisation hierarchy shown in Fig. 4 is only 1 and "opening" and "window" are said to be similar.

If the system fails to find structural correspondence between C_{cd} and any M_{cd} of the selected models, the system tries to bring out similarity hidden under the surface description of C_{cd} using knowledge stored in the DKB (*implied similarity*). The system transforms C_{cd} into an alternative description using knowledge in the DKB and repeats the similarity test using the notion of structural mapping. Heuristic rules used to transform C_{cd} are as follows.

1. Introduce objects which exist in M_{cd} but not in C_{cd} and add "part_of" relations between existing and new objects.
2. Add default properties if properties of objects in C_{cd} are missing. For example, if the occupancy capacity of a storage room is not given in C_{cd} then the default occupancy capacity, 0, is added to C_{cd}. Another example is, if there is a storey on the ground/first/second level, then the default height of the floor above the ground/first/second level is added to C_{cd}.
3. Infer relations between objects in C_{cd} and add them to C_{cd}. For example, if an object A is connected to an object B and B is connected to another object C, a new relation "A is connected to C" is added to C_{cd}. Another example is, if an object A is part of an object B and B is part of another object C, a new relation "A is part of C" is added to C_{cd}.

Once relevant models are found, the system compares the C_{sd} of the given case with the M_{sd} of retrieved models. If the solution described in C_{sd} is within the range of the solution described in M_{sd} of a model, then the case is said to satisfy the model. For a case to be determined to be *compliant* with the regulations, all "strong" models among retrieved models should be satisfied but "weak" models need not be satisfied.

4.5 Acquisition of New Regulatory Information

When the proposed solution C_{sd} does not satisfy any "strong" models retrieved as relevant, there is a possibility of relaxation. The decision on whether the

relaxation is to be granted will be made by a human expert. The system only provides relevant information stored within the system:

1. principles or issues behind relevant regulations, i.e. models at the upper levels of retrieved strong models which provide the reason why requirements represented as the retrieved strong models should be satisfied; and
2. previous similar cases, i.e. cases attached to retrieved models.

First, for example, if $(C_{id}, C_g, C_c, C_{cd}, C_{sd})$ of C90 shown in Sect. 4.2 is given to the system, M23 which represents Regulation D7 is retrieved as relevant to C90. The principles or issues behind Regulation D7 are represented as M6, M5 and M3 (see Fig. 2) and these three models are explained as follows.

> Principles behind D7 are as follows:
> the building ensures structural stability in case of fire
> therefore
> the building ensures safety in case of fire
> therefore
> the building ensures safety of people

Second, the system explains cases attached to models retrieved as relevant. In addition to the information about features of the building, argumentation recorded using an IBIS-like argumentation structure is also used to explain the key points in the debate surrounding the case. The argumentation structure can be easily converted into English-like text by following the links in the argumentation network. For example, the system explains argumentation on the proposal for relaxation suggested by the applicant in Fig. 3 as follows:

> The following proposal suggested by the applicant was granted: [the material used for the stair is timber and its fire resistance is one-hour]
> The proposal was supported by the following arguments:
> the applicant argued that: the proposal is a practicable solution when [the building is an existing building]
> the applicant argued that: under the circumstances in which [the building has a fire alarm system, smoke detectors and a tele nurse system] the proposal has the following consequences [the building ensures safety in case of fire]
> The proposal was opposed by the following arguments:
> the Local Building Control Authority argued that: the proposal constitutes non-compliance based on the facts that [the stair is a new stair]

The decision on whether relaxation is to be granted or refused is made by a human expert. Before asking the human expert to make a decision on relaxation, the system checks whether there is conflict between the proposed solution in the case and legal rules described in the same criterion in abstraction hierarchies of goals which is different from C_g. This test is needed because the regulations deal with various issues and it is very rare for any part of a building to serve only one purpose. Different standards may be given which use the same criterion to

resolve different issues. Therefore, when a criterion is evaluated for requirements in one issue, all requirements described using the same criterion with respect to other issues are checked and informed to the human expert.

After all relevant information—principles or issues behind the relevant regulations, similar cases and whether or not there is conflict with existing legal rules—is explained, the human expert is asked to make a decision on relaxation. If it is decided to grant relaxation, then the new rulings established in the relaxed case are integrated into the abstraction hierarchies. First, the system will hypothesize the category of buildings to which the new rulings should apply by looking at the M_c of "strong" sibling models which are similar to the case. Second, the system will collect "weak" sibling models which are similar to the case and generate the circumstances descriptions and the solutions by generalising the C_{cd}/C_{sd} of the case and the M_{cd}/M_{sd} of the collected similar "weak" models respectively. A new model will be created from the hypothesised category, the generalised circumstance description and the generalised solution. The strength will be set to "weak". This new model will replace the "weak" models whose circumstances descriptions and solutions are integrated into the new model.

5 Concluding Remarks

The Scottish Office records each appeal case in one file which contains documents, letters, drawings and photographs related to the case. 22 appeal cases were randomly selected in different subject areas and were formalised into 25 test cases. The prototype system has been implemented in Prolog using these cases. The regulations related to these 25 cases were also formalised into 102 models and four abstraction hierarchies were built. Background domain knowledge needed for processing 25 test cases was also gathered and formalised in the DKB. 25 test cases were supplied to the prototype system and the results were compared to the results by the human expert recorded in the case files. In 24 cases, the system retrieved models which represent the requirements against which the case appealed. In case 87, the applicant appealed against requirements of mechanical ventilation (represented as M112). Instead of providing some means of mechanical ventilation, the applicant insisted on using windows and doors as a means of ventilation and the system retrieved M108, M109 and M116, which represent the requirements of natural ventilation. Relaxation was granted in 17 test cases and 17 weak models were added to the MKB. There was no case among these 17 relaxed cases in which rulings could be combined with existing weak models (i.e. rulings from previous cases), since issues involved in the selected cases were very diverse and there were no cases which were similar to existing weak models.

A case-based approach as described in this paper has the following advantages over a traditional rule-based approach. First, it allows more flexibility in organising and maintaining legal rules than in a traditional rule-based approach. Although the abstraction hierarchies need to accommodate the representation and handling of exceptions (see Fig. 1 in Sect. 2), the interpretation hierarchies which exist among legal rules have been successfully represented as

abstraction hierarchies. Furthermore, since models are loosely connected in the abstraction hierarchies and inconsistency among models is tolerated (to a limited extent), new legal rules can be introduced more easily. Second, the representational framework described in this paper, i.e. the abstraction hierarchies together with the matching algorithm, has the following advantages:

- It allows the accessing mechanism to be more tolerant towards the representation of the input case by bridging terminological gaps between models and cases.
- It provides an efficient indexing mechanism by focusing on a particular group of models and pruning models unlikely to be relevant by using background domain knowledge before the similarity assessment begins. However, the use of background domain knowledge could be a drawback because it limits the generality of the system, and the system needs to provide extra facilities to acquire new background domain knowledge and to maintain consistency on the domain knowledge base.
- There is no need to compare the input case with all cases in the case library but only to compare models in the abstraction hierarchies to find out relevant legal rules, as the number of cases in the case library increases well above the number of models in the model knowledge base.

The system described in this paper does not make any decision on relaxation but absolutely relies on the human expert's discretion. It might be possible to extend the system to give advice on relaxation in two ways. First, it might be possible to project the input case to the argumentation structure of previous similar cases in order to predict the decision which the experts might make. A second possible approach is as follows. There is a set of factors related to each subject area, for example, in the subject of fire safety, there are several factors which usually are discussed such as: building type, occupant capacity, fire loading, number of storeys, fire-fighting access, alternative escape, etc. It might be possible to find typical patterns of argumentation, expressed in terms of those factors, in each subject area. Once these patterns of argumentation are formulated, it might be possible to apply this pattern to the input case and to predict the decision which the human expert might make.

6 Acknowledgements

We would like to thank the Scottish Office for allowing us to use appeal cases. Any views expressed in the paper about the building regulations are those of the authors alone.

References

1. Branting, L.K.: Building Explanations from Rules and Structureed Cases. International journal of man-machine studies **34** (1991) 797–837.

2. Conklin, J., Begeman, M.L.: gIBIS: A Hypertext Tool for Exploratory Policy Discussion. ACM Transactions on Office Information Systems **6:4** (October 1988) 303–331.
3. Collins, A.M., Quillian, M.R.: How to make a language user. In Organisation of Memory (1972). E. Tulving and W. Donaldson (Eds.). MIT Press.
4. Gentner, D.: Structure Mapping: A Theoretical Framework for Analogy. Cognitive Science **7** (1983) 155–170.
5. Holyoak, K.J., Thagard, P.: Analogical Mapping by Constraint Satisfaction. Cognitive Science **13** (1989) 295–355.
6. Kunz, W., Rittel, H.W.: Issues as Elements of Information Systems (1970). Working Paper. Center for Planning and Development Research, University of California, Berkeley, CA, USA.
7. Rissland, E.L.: Artificial Intelligence and Law. The Yale Law Journal **99** (1990) 1957–1981.
8. Rissland, E.L., Skalak, D.B.: CABARET: Rule Interpretation in a Hybrid Architecture. International journal of man-machine studies **34** (1991) 839–887.
9. The Building Standards (Scotland) Regulations 1981 (1981). Scottish Office.
10. Sergot, M.: Representing Legislation As Logic Programs (1985). Tech. Report. Department of Computing, Imperial College of Science and Technology, UK.
11. Sergot, M,, Sadri, F., Kowalski, R.A., Kriwaczek, F., Hammond, P., Cory, H.T.: The British Nationality Act as a Logic Program. Communications of the ACM **29:5** (May 1986) 370–386.
12. Stone, D.: Intelligent Information Systems for Building Standards. Proceedings of EuropIA 88, Paris (1988).
13. Susskind, R.: Expert Systems in Law (1987). Oxford University Press.
14. Thagard, P., Holyoak, K.J.: Why Indexing is the Wrong Way to Think about Analog Retrieval. Proceedings of DARPA Case-Based Reasoning Workshop (1989) 36–40.
15. Toulmin, S.: The uses of argument (1958). Cambridge University Press.
16. Walker, R.F., Oskamp, A., Schrickx, J.A., Opdorp, G.J.V., Berg, P.H.V.D.: PROLEXS: Creating Law and Order in a Heterogeneous Domain. International journal of man-machine studies **35** (1991) 35–67.
17. Waterman, D.A., Paul, J., Peterson, M.A.: Expert Systems for Legal Decision Making. Expert Systems **3:4** (1986) 212–226.
18. Yang, S.-A., Robertson, D., Lee, J.: KICS: A Knowledge-Intensive Case-Based Reasoning System for Statutory Building Regulations and Case Histories. Proceedings of the Fourth International Conference on AI and Law, Amsterdam, The Netherlands (June 1993) 254–263.

This article was processed using the LaTeX macro package with LLNCS style