



School of Informatics, University of Edinburgh

Centre for Intelligent Systems and their Applications

A Web-based Decision Support System for Divorce Lawyers

by

Scott Duguid, Lillian Edwards, John Kingston

Informatics Research Report EDI-INF-RR-0170

School of Informatics
<http://www.informatics.ed.ac.uk/>

November 2001

A Web-based Decision Support System for Divorce Lawyers

Scott Duguid, Lillian Edwards, John Kingston

Informatics Research Report EDI-INF-RR-0170

SCHOOL *of* INFORMATICS

Centre for Intelligent Systems and their Applications

November 2001

International Journal of Law, Computers & Technology, 15, 3, 2001, pp. 265-280.

Abstract :

A group of researchers from the EPSRC-sponsored Advanced Knowledge Technologies (AKT) project had independently developed four different ontologies covering the same domain: scientific knowledge management (academics, papers, conferences, etc). A reference ontology was developed to allow communication between these four ontologies. It was proposed that the reference ontology would benefit from incorporating the top level of the OntoClean ontology (Gangemi et al, 1998) because of the meta-properties that this ontology uses to support ontology structuring decisions. A merged ontology structure is proposed and the issues that arise in the merging process are analysed and discussed.

Keywords : Divorce law, matrimonial property, Java, Internet

Copyright © 2003 by The University of Edinburgh. All Rights Reserved

The authors and the University of Edinburgh retain the right to reproduce and publish this paper for non-commercial purposes.

Permission is granted for this report to be reproduced by others for non-commercial purposes as long as this copyright notice is reprinted in full in any reproduction. Applications to make other use of the material should be addressed in the first instance to Copyright Permissions, School of Informatics, The University of Edinburgh, 2 Buccleuch Place, Edinburgh EH8 9LW, Scotland.

A Web-based Decision Support System for Divorce Lawyers¹

SCOTT DUGUID, LILIAN EDWARDS and JOHN KINGSTON

ABSTRACT *This paper discusses the design and development of a prototype web based expert system for calculating financial provision awards when a couple divorce in Scotland. The function of the system is primarily to help trainee lawyers to understand the structure and content of the matrimonial property and divorce domain, and to access case law available on the web in a systematic fashion. The system currently focuses on evaluating what property belonging to the parties should be included in “matrimonial property”.*

We felt that an expert systems approach was appropriate for the domain of financial provision on divorce, since the modelling and application of expertise by an expert system is thought to be helpful in promoting “active learning”. An e-commerce style interface was used for gathering information, and the complete text of the Family Law (Scotland) 1985 is also available in the system, often with hyperlinks to the main narrative body of the system.

The system is designed to support an advisory, negotiating or mediatory role, rather than an adversarial one, since Scottish solicitors have a preference for seeking consensual negotiated settlements. A practical and useful future extension to the system would be to permit the drafting of electronic documents resulting from negotiated settlements, such as minutes or joint minutes of agreement.

Introduction

This paper discusses the design and development of a prototype web based expert system for calculating financial provision awards when a couple divorce in Scotland. Issues of financial provision (ie allocation of matrimonial property and associated issues) arise in almost every divorce case, although in many cases they are settled by informal or formal agreements rather than litigated to a judicial conclusion. Indeed, ‘it is the financial aspects of divorce together with child care arrangements which are overwhelmingly the most crucial and disputed matters in modern divorce’.² Although the splitting of assets on divorce is often portrayed in the media as highly adversarial, in fact Scottish solicitors

Correspondence: Scott Duguid is a Research Assistant in the Faculty of Law, Edinburgh University; Lilian Edwards is Senior Lecturer in Law, Edinburgh University; John Kingston is at the Artificial Intelligence Applications Institute, University of Edinburgh.

(unusually in the world perhaps) have a preference for seeking consensual negotiated settlements and such agreements are formalized as minutes or joint minutes of agreement in around a quarter of all divorces, with far more informal compromises being rubber stamped by courts at the end of the day.³ The legal profession thus has much to gain from a system that potentially improves the post law degree education of trainee divorce lawyers, particularly as the body of case law interpreting the 1985 Act is already large and continuously expanding. Moreover, in a climate where most Scottish law firms now have on-line access and Scottish legal cases are increasingly available electronically, and often on-line and for free, a web-based system with access to relevant cases in digital form, and possibly other electronic documents such as styles of minutes of agreements, has obvious potential.

The function of the system described in this paper is thus primarily to help trainee lawyers who are either new to, or rusty on, the nuts and bolts of the matrimonial property and divorce domain, both to understand the structure and content of the domain, and to access case law available on the web in a systematic fashion. Moreover the system is designed to support an advisory, negotiating or mediatory role, rather than an adversarial one. The thesis advanced is that an expert system model is more appropriate to training in this kind of domain than classic question-and-answer type systems, since it resembles and automates the type of aid given by an experienced lawyer to a novice one in the off-line world for centuries.

In terms of the technology used, we felt that an expert systems approach was appropriate for the domain of financial provision on divorce. Although the document base in the domain (ie cases and statutes) is not enormous, it is sufficiently large that access to all and only the relevant documents is desirable. This can be achieved by a combination of structured hypertext documents and an expert systems approach. Tasks required in the domain of financial provision, such as information gathering and storage; arithmetical calculations; and applications of rules, may suitably be performed by an expert system. In addition to this, the domain is not one dominated by common sense (notoriously hard to implement in a computerized system), but complex enough that expertise is required and not one simply devoted to number-crunching or document retrieval.

The Domain: the 1985 Act, 'Matrimonial Property' and Fair Sharing

In modern Scots family law, issues of financial provision on divorce are governed by the *Family Law (Scotland) Act 1985*.⁴ References below are to the 1985 Act unless stated otherwise. Although periodical allowance orders remain available, the principal philosophy underlying the Act is that divorce should be as far as possible a 'clean break' between the parties.⁵ Section 8 of this Act provides that the court has discretion on divorce to make a package of orders relating to financial provision, including capital sum orders, property transfer orders, periodical allowance orders and a number of incidental orders. How this discretion should be exercised is crucially guided by the five principles found in s 9. The key s 9 principle is s 9(1)(a), which states that 'the net value of the matrimonial property should be shared fairly between the parties to the marriage'. 'Matrimonial property' is defined in s 10(4) as:

... all the property belonging to the parties or either of them at the relevant date which was acquired by them or him (otherwise than by way of gift or succession from a third party):

- (a) before the marriage for use by them as a family home or as furniture or furnishings for such home; or
- (b) during the marriage but before the relevant date.

What constitutes 'fair sharing' is defined in s 10(1), which states that:

In applying the principle set out in section 9(1)(a) of this Act, the net value of the matrimonial property shall be taken to be shared fairly between the parties to the marriage when it is shared equally or in such other proportions as are justified by special circumstances.

The 'relevant date' referred to in s 10(4) acts as a cut-off point for the valuation and assessment of matrimonial property, and according to s 10(3) two dates may qualify: '(a) ... the date on which the parties ceased to cohabit'; and '(b) the date of service of the summons in the action for divorce'. Items that fall into the 'matrimonial property' are valued as at the relevant date.

Because equal sharing may be inequitable in a number of 'special circumstances', and indeed because parties may no longer own all the assets that once comprised matrimonial property at the actual date of divorce, the courts are granted considerable discretion to distribute the matrimonial property as it exists at the date of divorce in unequal proportions, and indeed to make such awards in general as are (a) justified by the s 9 principles and (b) reasonable having regard to the resources of the parties (s 8(2)). 'Special circumstances', as non-exhaustively defined in s 10(6)(a-e) include (a) the terms of any agreement; (b) 'the source of funds or assets'; (c) destruction or dissipation of property; (d) nature and use of the matrimonial property; and (e) liability for expenses of valuation or transfer.

In addition to s 9(1)(a), the four remaining s 9 principles deal respectively with: s 9(1)(b), economic advantage or disadvantage; s 9(1)(c), economic burden of childcare; s 9(1)(d), dependence by one party on financial support of other party for a limited period; and s 9(1)(e), potentially unlimited support based on financial hardship. As may be seen, these are considerably more discretionary than the basic structure of s 9(1)(a) and are likely to require the use of a different paradigm for implementation than the rule based reasoning adopted in the s 9(1)(a) module to date.

The Expert System: User, Aims and Function

User

The current tool is intended as a decision support system in the domain of financial provision on divorce in Scots family law. Its purpose is first, to facilitate learning in the domain, and second, to act as an advisory system/*aide memoire* after initial training has been undertaken. The current intended user is the trainee divorce lawyer. It is anticipated that the user will have prior general legal training, but residual or incomplete knowledge of the specific domain modelled in this system, ie the relevant sections of the Family Law (Scotland) Act 1985 and the surrounding case and statute law.

Aims

Legal expert systems have in the past been developed to operate both in advisory and adversarial roles, the former predominating over the latter. Capper and Susskind's Latent Damage System is a well-known example of the former approach; Ashley's HYPO is

arguably an example of the latter. They have not however often been conceived of as teaching and learning tools. Traditional legal computer programs aimed at assisting students in law schools, such as the CALI, TLTP and IOLIS projects of the 1990s, tended predominantly to be based around what were essentially question-and-answer models. Many ingenious approaches were used to disguise this fact: learning systems were developed based around Socratic dialogues, textbook chapters, real-world problem based tutorials or fact patterns, flowcharts, concept graphs, or at the most basic level, yes/no questions and multiple choice quizzes. But all of these models, however sophisticated, effectively ask the user/student to recall legal information with or without prompt, which is then checked as correct or incorrect, and then perhaps to manipulate it in some way that demonstrates more than sheer recall, eg synthesis or connection of concepts.⁶

Expert systems represent a very different kind of intellectual activity. Instead of the user being asked to retrieve data from memory, or in more advanced systems, to analyze that data once retrieved, the user is typically shown in detail how a program written with the aid of a domain expert—modelling the expertise of a skilled user—would reach a desired goal given the known factual inputs to the problem domain. The thesis that is advanced in this project (and in previous work undertaken by Edwards⁷ is that this kind of informed participation in problem solving, although apparently more passive than ‘classic’ computer assisted problem solving, actually is more helpful in promoting what is slightly optimistically termed ‘active learning’ than the more traditional approaches. Learning at the elbow of an expert—the ‘apprenticeship’ model—has hundreds of years of success behind it; learning unaided except by question prompts, and memory work, tends only to promote a mechanical kind of learning more appropriate to arithmetic than law. In Edwards’ previous work in the Scottish succession domain,⁸ students were asked to solve a tutorial problem which required them to calculate the division of an intestate estate using the SUCC expert system, which was made available over the Law network. Results were not statistically analyzed, given the relatively small number of students who took up the challenge, but the anecdotal outcomes were that performance both in subsequent real-world tutorial groups, and in questions of the same ilk in the exam, significantly improved. Students also reported finding the system more pleasurable and interesting to use than traditional question-and-answer tutorials prepared in the same domain as part of the TLTP project (and in one memorable case, more pleasurable—and logical—than a real human-led tutorial).

Function

As already noted, expert systems can be used in an adversarial or advisory capacity. The ‘adversarial’ model is however often inappropriate to divorce cases where negotiation and settlement are often seen as preferable to an adversarial court-based approach, on grounds of speed, cost and lack of emotional trauma. In Scotland at least, ‘defended family actions are the exception rather than the rule and much if not most of a family lawyer’s time will be taken up with negotiation, settlement, and drafting of agreement’.⁹ Divorce settlements are also increasingly worked out with the aid of mediators and solicitor–mediators. The advice provided by an appropriate legal expert system, as interpreted by a lawyer, may help facilitate ‘bargaining in the shadow of the law’, where negotiations are informed by some idea as to the orders for financial provision a court might be likely to make in an actual disputed case. The functions we expect our system to support then are mainly the advisory and negotiation roles undertaken by divorce lawyers. A practical and useful future extension to the system would be to permit the drafting of electronic documents resulting

from such negotiation, such as minutes or joint minutes of agreement, and we hope to seek further funding to add this functionality in time.

Reasoning and Representation in the Domain

As discussed above, a substantial and important part of the reasoning process in a divorce case is to evaluate what property belonging to the parties should be included in 'matrimonial property'. As previously stated, s 10(4) of the Act provides a definition of what constitutes matrimonial property. For any item of property, *if*:

- it was acquired 'before the marriage for use by them as a family home or as furniture or plenishings for such home' or 'during the marriage but before the relevant date';
and
- it is owned by either or both of the parties; *and*
- it was not acquired 'by way of gift or succession from a third party'

then its value at the relevant date is placed into the matrimonial property 'pot'. The above section clearly lends itself to formalization in rules. Problems may still arise as to unclear facts or ambiguity in language but some of these may be dealt with by direct reference to the full text of case law. For example, there may be, as in *Buczynska v. Buczynski* 1989 SLT 558, a dispute over when the relevant date occurred. A link to this case will help to elucidate the open textured concept of 'cessation of cohabitation'.

However, not all parts of the financial provision domain are so clearly rule based. Judges have considerable discretion in deciding when to share matrimonial property unequally where there are deemed to be 'special circumstances', and may in addition apply any of the principles in s 9(b-e). Finally, the award they make is determined not just by the s 9 principles but with regard to the total resources of the parties and may be constituted with great flexibility using all the orders available in s 8.

The Scottish financial provision domain may thus be seen as an interesting hybrid between a rule-based, relatively predictable, domain, of a type similar to those found in civilian matrimonial property regimes, and highly discretionary divorce domains such as are typically found in England, the US common law states and Australia.

Discretionary reasoning is typically difficult to model using conventional expert systems technology, ie rule based reasoning. In their nature rules tend to produce yes/no results; while discretionary reasoning tends to identify a subset of the most likely or desirable results out of the set of possible results. Some rule based approaches do exist to deal with this issue, eg the use of fuzzy logic; however these were not suitable to the project at hand where the aim was to use cheap technologies suitable to (a) rapid modelling and (b) implementation via the Web.

The problems of dealing with a highly discretionary financial provision domain have already been addressed by Stranieri and Zeleznikow in their *Split Up* system, a system for predicting judicial decisions under the Family Law Act (1975) of Australia. The 1975 Act is described as a discretionary domain because 'it makes explicit a number of factors that must be taken into account by a judge in altering the property interests of parties to a marriage, but the statute is silent on the relative importance of each factor. Different judges may, and do, reach different conclusions, even when they agree on facts, because each judge assigns different relative weights to factors.'¹⁰

The builders of *Split Up* found that rule-based reasoning was only of limited value in a discretionary domain. *Split Up*'s approach to solving the problem is to employ a neural

network to determine the percentage split of property in a divorce case, based on a large bank of previous routine cases. Note that this is distinct from case-based reasoning, also rejected by the builders of *Split Up*, where 'a select number of leading cases' are used. Case-based reasoning was rejected as a reasoning paradigm because it tends to model a small number of leading superior court cases that embody novel legal arguments. When trying to determine the distribution of matrimonial assets, what is looked for is the routine distribution in very many cases; not the exceptional, stand out cases where some quirk of the law makes the distribution uncertain or unprecedented. Neural networks, which essentially produce results derived statistically from a very large number of typical cases in a domain are ideal for this sort of routinized domain.

A perceived problem with neural networks in legal expert systems however is the lack of facilities for providing explanations. They are often evocatively described as a 'black box'. The builders of *Split Up* have at current prototype stage addressed this problem by the use of Toulmin arguments: essentially once the percentage split of property has been generated from the neural network, the facts of the case are separately processed using logic to generate post-rationalized arguments that *may* explain the property division the NN has generated. There is a logical connection, but no actual causal connection, between the result generated by the neural net and the arguments(s) that back it.

Since the system being described in this paper is currently in the prototype stage, the problems of the discretionary aspects of the domain have as yet not been fully addressed. It may be however that they will create no enormous difficulties. The reason for this is that there are no immediate plans for the system to generate an exact percentage split of property. Rather, the objective of the system is first, to calculate what constitutes matrimonial property; and then secondly, to generate *arguments* based on the facts as applied to the discretionary elements of the s 9 principles, to see if there are reasons to diverge from the norm of equal sharing (and if so, to what extent). Arguments, both for and against divergence, will be backed up by hyper-linked reference to appropriate case law.

There were a number of reasons why we decided not to attempt to generate a discretionary result rather than merely allow arguments to be asserted. Given that the system is primarily aimed at training novice lawyers rather than actually advising clients, the final exact result is not the goal of the system. Rather the aim is to transfer to the novice the processes an expert goes through in reaching that 'guesstimate'. In terms of learning outcomes, it is more useful for a trainee lawyer to have to consider if case law on s 9(1)(b) (say) is favourable or unfavourable to one spouse given the facts of the case under study, than it is for him to be presented as a *fait accompli* with a percentage split worked out by a neural network or similar technology. Even if we had felt motivated to provide a final exact result (perhaps as an 'advisory' alternative to the 'training' mode) it is doubtful if we could have obtained appropriate data on enough routine cases to make a neural network or other statistical solution viable. Scotland is only a small jurisdiction and as previously noted, many cases are settled rather than litigated. Even if we had used all available public court records of financial provision on divorce cases, we would still have run the risk of deriving results from a pathological rather than a routine sample of cases.

The problem of how to model and derive results from a discretionary legal domain remains an interesting technical and theoretical one that we hope to pursue in later stages of this project. Concentrating on the training and advisory goals however, it will be interesting to see at the assessment stage how far typical users feel the system is flawed by not providing exact percentage splits after the first rule-based module of the system is completed.

Web-based Access to Legal Documents

Electronic creation and storage of legal documents has evident benefits for such a manifestly textual domain as law. Information retrieval systems, such as Lexis and Westlaw have proved invaluable resources for lawyers. Increasingly, too, documents are being placed on the Web for free and without proprietary restrictions of copyright. Both the British and Irish Legal Information Institute—or Bailli (at <http://www.bailli.org>)—and the Scottish Courts Website (<http://www.scotcourts.gov.uk>) provide searchable on-line access to cases.

Convenient as information retrieval databases are, they have particular flaws. Many on-line resources only have a finite number of cases on-line. The Scottish Courts Website only posts cases from the Court of Session filed since 1998, for example. Moreover, the full potential of hypertext is rarely exploited in databases, where the documents themselves are frequently structured largely in plain text.

A more fundamental issue regards the reliable retrieval of relevant documents to a case. As Susskind has noted, 'We are a long way from the ideal, that of these systems [information retrieval systems] having total and precise recall, being able, that is to say, to retrieve *all but only* the relevant documents for a user's particular purposes'.¹¹ For a range of reasons, search by keyword can both bring up both too many irrelevant documents and at the same time can miss certain relevant ones. This problem often exists even where the user is relatively adept with search strategies.

When working in a relatively well-defined domain, expert systems can be of great use in guiding users through a set of questions that are necessary and sufficient for their needs. On the other hand, the advice of an expert system is merely a starting point for the legally trained user, who must employ discretion in interpreting the advice: 'it would be ... unacceptable if the user of an expert system in law did not look beyond, but always accepted without further query, the advice offered during a consultation with his system'.¹² As much access as possible to primary case sources is necessary if the system is to be both transparent and self-justificatory. Ideally, the user should have on-line access to the full text of a source cited, as well perhaps as to an abstract thereof (whereby users could verify the relevance of the case before consulting it in full). Support of on-line access to relevant cited cases is one of the main goals of the current project.

The User Interface: Information Gathering and Learning the Domain

Because the purpose of the system is as much to facilitate learning as to provide useful advice, emphasis has been placed on the user interface and especially on full explanations for any information requested or advice offered. Zeleznikow and Hunter have noted that '[l]earning by doing has long been seen to be more effective than passive learning. Legal expert systems clearly have their place in legal teaching, even where they are designed for other purposes. Legal expert systems which are specifically designed to teach an area of law are even more useful for pedagogical purposes'.¹³ The information-gathering process is therefore central to our methodology.

Essentially, the problem of how the system helps the user learn can be broken down into three sub-problems: how the system gathers information on a user's case (input); the format of the advice/help provided (output); and the procedures that govern how the input and output is ordered. Again, to a certain extent these problems depend on each other. Let us take the example of the calculation of the relevant date. To establish the relevant date and the date of marriage, three separate dates (the date of marriage, the date of cessation of




Figure 1.

ohabitation, and the date of the summons) must be provided. In order to do this, the user will have to be provided not only with a definition of all three of these terms (as relating to the s 10(3) of the 1985 Act), but also will need to know about related concepts, such as that of matrimonial property. The user may also require advice if there is a dispute as to the relevant date. In turn, future advice provided by the system on other issues will be affected by the established relevant date.

An initial prototype for the current project was an undergraduate project, written at the University of Edinburgh,¹⁴ which attempted to tackle the same domain. This project essentially consisted of a series of diagrammatically represented procedures (Figure 1).

The procedures were followed through a 'question-answer' information-gathering session. The diagrams were written in *ProForma*, a graphical language first developed for knowledge acquisition in medical domains. Each node on the diagram represents a task, which can be condensed into a series of sub-tasks. For several reasons this approach was inappropriate for our needs. First, although the essential procedural structure was sound, it was insufficiently tailored to the user's needs. Occasionally information was requested in an order that was far from intuitive. Second, while the domain can be modelled diagrammatically, such a representation only shows that there is a series of procedures—it does not explain the statutory reasons for the procedure. Legal knowledge is highly textual, to the extent that 'text has been absolutely fundamental to legal practice and to the administration of justice'.¹⁵ This suggests that a narrative approach to legal teaching is preferable to a diagrammatic one. However, this is not to say that diagrams cannot provide a good conceptual model of legal reasoning.

Despite these reservations, the previous prototype did provide a basic essential set of procedures for the system, albeit one that had to be modified. If the user is to learn effectively, information should be supplied and requested in a manner that is both (a) intuitive to the user, and (b) logical in respect of the statute. For example, it would make little sense to allow a user to use a module dealing with 'special circumstances' (as defined in s 10(1) and s 10(6)) before the resources of the parties had been assessed, and the 'matrimonial property' evaluated. Once the sequence of modules has been established, there should be a method of requesting information that is consistent with the nature of the knowledge required. Rule-based legal expert systems have often tended in the past to use a backward-chaining 'question-answer' methodology,¹⁶ where an answer to one question leads to the asking of another in a 'tree-like' manner until a top level goal is proved true or false. This is particularly appropriate in many legal domains which are inherently 'single goal oriented' (eg 'can jurisdiction be founded in a particular Scottish sheriff court?'; 'has an action for latent damage been time-barred?').

The problem with such an approach in the current domain is that it is less 'single goal oriented', and more a domain where the information supplied initially may lead to a number of less or more satisfactory goal states. Similar domains would include tax planning, or legal protection of an industrial product. Rather than establishing the existence or absence of a particular state of affairs, the current domain requires that a list of all resources of the parties be entered into the system and the consequences under the 1985 Act followed through from there. Entering the resources sequentially is not only likely to be tedious for the user, but also emphasizes the information gathering process at the expense of explaining the statute. The solution adopted to this problem is for the user to input all the relevant resources into an inventory at one go. This model we have described as an 'e-commerce' interface, because consumers have become accustomed over the last few years to on-line forms as found on web sites such as Amazon.com, Yahoo!, Ebay, etc. The user can see all the information represented on one form, rather than distributed through successive screens and can make amendments before finally submitting it. Moreover, a form-based approach can aid learning, as the structure of the form can represent visually to the user the hierarchical structure of the domain knowledge. Susskind has warned against the 'input of problem data at one fell swoop at the start of the dialogue',¹⁷ but the inventory of resources is only one constituent part of the information-gathering process in the domain of financial provision. The system is still interactive in the sense that the advice given and questions put to the user later will be contingent on the knowledge provided on the inventory.

Implementation

Development Tools

There were three essential requirements in terms of development tools. First, we required a language that was object-oriented that could model the structure of knowledge required in the domain. Second, we needed a rule-based inference engine to perform the legal reasoning. Last, we needed a method of building a user interface that was compatible with the worldwide web, specifically HTML forms. The programming language we used was Java, which not only has the object-oriented capabilities we required, but also supports a variety of Web-related technologies. The inference engine we used was JESS (Java Expert System Shell), a rule-based language that started life as a Java implemented extension of CLIPS,¹⁸ but increasingly has begun to support various methods and classes for more efficient interaction with Java code. A JESS rule engine can be embedded into and manipulated by Java code as a 'Rete object'. In addition, objects in Java can be replicated in the rule engine as JESS objects. When a Java object is manipulated, its counterpart in JESS changes its properties also and vice versa. This brings the convenience that procedural and rule-based tasks can be separately conducted in the appropriate language.

There were several options regarding the user interface. Client-side approaches such as Java applets and Javascript were rejected. Javascript handles HTML forms, but has limited processing power, while applets are not good for handling HTML and, moreover, are better suited to a 'question-answer' style approach than a form-based one. Instead, a server-side model seemed more appropriate. The interface was built using Java Server Pages (JSPs). JSPs essentially look like HTML pages, but contain Java code that produces dynamic HTML content. They work in a similar way to servlets (and JSPs are often

used in conjunction with servlets), but are better for producing Web pages with a lot of variable content. Java objects (including a JESS inference engine) on the server side can be accessed by the JSPs, which make them useful for tracking a session over a series of pages. They are also potentially useful for storing data on the server side, which would allow a user to complete a run of the system over a series of sessions if required, or access multiple cases as and when required (although such an implementation would bring up security issues).

Essentially, we have applied an 'e-commerce' architecture to a legal expert system. Another potential programming language we could have used was Perl, in conjunction with CAPE,¹⁹ a rule-based language compatible with Perl. Perl does support server-side processing, but we felt that JSPs provided features for more rapid development of dynamic content.

Access to Documents

The complete text of the Family Law (Scotland) 1985 is available in the system. It can be perused at leisure, but is frequently hyper-linked to in the main narrative body of the system. Citations of cases will ordinarily be hyper-linked to abstracts of the case. The user can use the abstract to confirm the relevance of the document in question. Where appropriate, the abstract will be followed by a relevant extract from the ruling proper. For many abstracts, it is hoped there will be access to the full case also, either through an external Web link, or linked to an internally stored copy of the document (where copyright permits). Abstracts and internally stored full documents will both provide hyperlinks to referenced cases and statutes.

Implementing the User Interface

In trying to build a system that was user-friendly, we stuck to a few guiding principles. Simplicity was important, and this was achieved by keeping the information-gathering process as intuitive, practical and uncluttered as possible. This was desirable since the legal user may not necessarily have good computer skills. Frequently a balance had to be struck between keeping the interface simple, and requesting the right complexity of information. In the initial prototype, we have tended to sacrifice more complicated representation of knowledge for a more user-friendly interface when conflicts have arisen. To keep the interface uncluttered and clear, we felt that the narrative sections of the system had to be separate from the information-gathering forms, while maintaining a close relationship between the two. This was achieved by using juxtaposed frames, one containing the form, and the other with narrative explanations of the statute and other information. Fields on the forms were hyperlinked to explanations in the narrative frame.

Figure 2 shows a typical user interface screen. The left frame contains the form for information gathering, which must be completed before going on to the next stage. The form is hyperlinked directly to the right frame, so the user can access information on how to fill in any field, with an explanation of its relevance to the statute. In addition, the right frame provides a broad narrative explanation of the statute as it pertains to the particular form in question. Navigation in the right frame can be pursued *ad hoc*, as hyperlinks permit the return to the main narrative at any time.

Figure 2.

Implementing the Procedural Flow

The procedural flow of the system, based on and adapted from previous diagrams in Dale, is implemented as a series of modules. The first module (to date the only one implemented) represents s 9(1)(a). It consists of four sub-modules. It gathers information and compiles a comprehensive inventory both of 'matrimonial property' and overall assets not qualifying as matrimonial property owned by the parties (eg a house inherited from an aunt). The module as so far implemented excludes the discretionary issue of 'special circumstances' so that at the end of the module, each spouse is presumptively allotted one half of the net value of the matrimonial property. 'Special circumstances' will be completed during the remainder of the project, as will modules relating to ss 9(1)(b)–(e). If time allows, a final module will provide insight into the 'package' of orders that the court might make in the case to implement the division of assets, and if possible, an extra, document drafting module will be added to generate an appropriate minute of agreement.

Sub-module 1A requests general information about the parties relevant to the case. Sub-module 1B seeks to establish the date of marriage and relevant date. Sub-module 1C asks for a complete inventory of the resources belonging to the parties, and sub-module 1D asks about any debts (in order to establish the *net* matrimonial property). The whole module 1 results in a complete list of resources, split up into matrimonial and non-matrimonial property, and with debts deducted, with explanations for each decision available.

Sub-module 1A asks for general details about the parties, such as their names, ages, number of children (in total and those under 16), current profession and annual income, and current residence. These are questions that may only be relevant to points of law arising later according to the layout of the 1985 Act but which are intuitive to ask first, modelling the user interface on a typical solicitor–client consultation. A deficiency in the initial prototype was that the user was immediately prompted for esoteric information relating to the statute, whereas we anticipate it being used by lawyers who may not be

Decision support for divorce lawyers Microsoft Internet Explorer

File Edit View Favorites Tools Help

Address http://kramer:8080/divorce/Indeframe.jsp Go

Links Customize Links Free HTML Windows RealPlayer Decision support for divorce lawyers

Inventory of Resources

Home Add new

Type of resource	Value (£)	Acquired When?	Owner	Acquired by gift or succession?	Matr. prop?	
<u>Home:</u> 42 Thirlstone Rd	£ 160 000	During marriage	Both jointly	<input type="checkbox"/>	?	Delete
<u>Home contents</u>	£ 0		Both Jointly	X	✓	?
<u>Cash (e.g. from bank):</u> Bank of Scotland	£ 20 000	During marriage	Husband	<input type="checkbox"/>	?	Delete
<u>Cash (e.g. from bank):</u> Inheritance	£ 15 000	During marriage	Wife	<input checked="" type="checkbox"/>	?	Delete
<u>Pension</u>	Click to add details of pension Go					Delete
<u>Other</u> Painting	£ 50 000	During marriage	Husband	<input checked="" type="checkbox"/>	?	Delete

Submit

Error on page.

Start Paper... Mouse... Conn... Test... Hobo... Build... Tomc... Decis... Local intranet 16:42

The inventory of resources

We know the relevant date of the marriage. Now the matrimonial property value should be evaluated. On this form you will enter the resources and assets of the parties. You should enter:

1. Any resource or asset that you believe might be considered matrimonial property as defined in s.10(4) of the Act
2. Any other resources belonging to either or both of the parties

Why (2)? Although any such resource will not be included in the

Last modified: 1 Apr 1992
Released: 1 Apr 1996

Figure 3.

familiar with the 1985 Act. Incomes, for example, become relevant only when addressing principles 9(1)(b) onwards and are irrelevant to s 9(1)(a) but are useful to gather at an early stage. Some facts also exclude irrelevant parts of the consultation at a later stage. If there are no children under 16, for example, there would be no need for the system to take into account s 9(1)(c) in a later module. If there *are* children under 16, it is appropriate to defer ‘deeper grain’ questions to the relevant module.

Sub-module 1B establishes the relevant date. This sub-module seems relatively straightforward, asking as it does for only three dates: the date of the marriage, the date of cessation of cohabitation, and the date of the delivery of the summons. However, there is a considerable amount about the statute that the user needs to know before this information may be reliably provided. Not only do terms such as ‘cessation of cohabitation’ need to be explained, but the term ‘relevant date’ itself depends on the definition of matrimonial property. This illustrates the necessity of having a narrative explanation alongside the information-gathering process.

A second, more complicated issue arises around the relevant date. What about those cases, such as the aforementioned *Buczynska v. Buczynski* 1989 SLT 558, where there is a dispute over the relevant date, ie the date when cohabitation ceased. As in that case, the value of the net matrimonial property may vary widely from one date to the other. One possible solution would be to allow the user to provide two potential dates for the cessation of cohabitation. On that basis, two potential net totals for matrimonial property could be calculated and printed out to the user. Such a ‘what if?’ capability would be useful and relatively easy to implement.

Sub-module 1C was the most complicated part of the s 9(1)(a) module to implement. It consists of an 'inventory of resources' or table containing items of property and their attributes. The table can list anything from one item to a potentially limitless number. Items of property were classified in types for ease of user data entry, although these types are not, it should be noticed, official statutory definitions. The types selected to date are *Home* (to include other items of heritage in the guidance supplied), *Home contents*, *Cash and investments*, *Pensions and life assurance policies*, and *Other*. Guidance is provided under *Other* on 'unusual' assets such as 'foreseeable resources'—damages claims, redundancy payments and the like. Each item must be provided with a value. For all items other than the home contents this must be greater than 0 or an error will result. If the item is to be included in the matrimonial property, this value will be included in the gross value of that property; otherwise, the value of the property is considered to be part of the owner's own resources. The final three attributes gather information sufficient to decide if the item falls within matrimonial property according to s 10(4) of the Act: namely, time asset was acquired (before, during or after the marriage), ownership of the asset (husband, wife, both or neither), and how acquired ('by gift or succession', or not—a Boolean value, which if false excludes the asset from 'matrimonial property').

The final three attributes do not apply to 'Home contents', since the system presumes, following s 25 of the 1985 Act, that in the absence of other evidence all home contents are owned jointly and are matrimonial property. Any disputed household items must be explicitly entered as an *Other* item. For two further asset types, the 'Home' and the 'Pension', which are likely to be the most important assets in many marriages, further information is required. For this reason, for each 'Home' and 'Pension' item, the user must go to a separate form. In Figure 3, the user has filled in the other form (with the relevant fields being retained on the inventory proper), but has not yet filled in the relevant form for the pension. The 'Go' button will take the user to that form.

When the user has filled in the inventory to his/her satisfaction, it is submitted, checked for errors and the user is presented with a list of resources with an indication of which ones constitute matrimonial property. This list can be either accepted or rejected by the user. The user is then led to a similar inventory relating to debts in sub-module D.

Sub-module D inventories debts relating to assets. Debts are removed from the gross total to form the net matrimonial property. Heritable debts must be taken off the relevant item of heritage. A definition of a debt for this purpose is provided in the statute under s 10(2).

The Future

As can be seen from the above, much work remains to be done before the prototype is complete. When this is achieved, it is hoped to have the system assessed by practising divorce lawyers and, after appropriate modification, to trial it in a number of law firms where feedback and comments can be gathered on its use as a training and advisory system. We will then seek funding to turn the prototype into a fully functional and maintainable system and to undertake both commercial and academic dissemination. We also hope if possible to see what use if any can be made of such a system in the mediation environment and if it is appropriate for lay mediators as well as solicitor-mediators.

Turning more generally to the future, it is noticeable that despite the exponential growth since the late 1990s in primary law sources on-line, especially on the free-to-all Web, a similar growth in on-line legal guidance systems, or secondary sources, utilizing this

primary source material, has not yet taken place, despite much futurology forecasting such developments.²⁰ Much has been written in the last few years about digital democracy and access to justice, yet as long as on-line law is restricted to primary legal sources without further explication, it is effectively of little use either to the lay public or indeed to the legal profession, whom surveys have shown have little time to absorb new primary legal sources without guidance or training. Such on-line legal advice systems as have so far been developed have tended to be confined to private legal firm intranets built at great expense and jealously guarded to provide competitive advantage, or systems provided to the public but effectively priced to exclude all but high-paying corporate clients.²¹ At the other end of the scale, some low cost legal calculation or document drafting packages have been developed, often marketed by non-law organizations direct to the public and cutting out the legal profession entirely.²² However such systems do nothing to explain the law to lay people (hence, entirely failing to educate or empower), nor to improve training standards in the legal profession; and anecdotal evidence suggests that most lay people still feel more comfortable being represented by a lawyer than acting for themselves in the legal domain.

The current system is aimed at an intermediate goal of creating a potentially fairly cheap on-line interactive training or 'first-go' system, using free Web-based sources, which may improve both the standard and the speed of legal advice given at the 'routinized' end of the legal market mostly commonly encountered by non-corporate clients. Family law is one of the areas of law most commonly encountered by the private legal client and yet it is an area where standards of legal advice and attention may conceivably be very variable due to lack of high profit margin and the consequent downgrading of the work and lack of time for training. One of the goals of the current project is to demonstrate that free on-line primary legal sources on the Web can be combined with software built using open-source tools to provide a benefit to small- and medium-sized law firms, as well as to the general public in terms of improvement of the standard of legal services delivered in a mass market field such as matrimonial property work. We hope to assess more fully the utility and potential market for the prototype once the implementation and assessment phases have been completed in 2002.

Notes and References

- 1 The work described within was funded by a one year ESRC grant and additionally by the Moray Fund administered by Edinburgh University. Our thanks to both institutions.
- 2 L Edwards and A Griffiths *Family Law* W. Green, Edinburgh, 1997, p 340.
- 3 *Ibid*, ch 15.
- 4 See A Hall-Dick *Family Law (Scotland) Act 1985* W. Green: Edinburgh, 2000.
- 5 *Ibid*, p 341.
- 6 This is not to say that CAL programs have not been developed as other than expert systems which are not based around the information-retrieval paradigm. Legal 'community' games modelled after popular games such as *The Sims* are models that fit neither camp.
- 7 L Edwards and J Huntley 'Creating a civil jurisdiction adviser' *Law, Computers and Artificial Intelligence* Vol 1, No 1, pp 5–40, 1992; L Edwards 'Building an intestate succession adviser: Compartmentalisation and creativity in decision support systems' in Z Bankowski, I White and U Hahn (eds) *Informatics and the Foundations of Legal Reasoning* Kluwer, Dordrecht, 1995, pp 311–323.
- 8 Edwards, *op cit*, note 7.
- 9 Edwards and Griffiths, *op cit*, note 2, p vii.

- 10 Stranieri, J Zeleznikow, Gawler and Lewis 'A hybrid rule—neural approach for the automation of legal reasoning in the discretionary domain of family law in Australia' *Artificial Intelligence and Law*, Vol 7, pp 153–183, 1999 (see pp 153–154).
- 11 R Susskind *The Future of Law* Clarendon Press, Oxford, 1996, p 109.
- 12 R Susskind *Expert Systems in Law* Clarendon Press, Oxford, 1989, p 62.
- 13 J Zeleznikow and D Hunter *Building Intelligent Legal Information Systems* Kluwer, Boston, 1994, p 71.
- 14 M Dale 'Building knowledge based decision support for divorce lawyers on the internet' MSc Project report, University of Edinburgh, 1999.
- 15 Susskind, *op cit*, note 13, p 107.
- 16 See for example the work done in the field of civil jurisdiction by Edwards and Huntley, *op cit*, note 9, and Susskind on the latent damage system (R Susskind *Transforming the Law* Clarendon Press, Oxford, 2000).
- 17 Susskind, *op cit*, note 14, p 66.
- 18 See www.ghgcorp.com/clips/CLIPS.html or JC Giarratano, *Expert Systems: principles and programming* 3rd edn. Brooks-Cole, 1998.
- 19 See <http://www.hcrc.ed.ac.uk/~robert/CAPE/>.
- 20 See especially Susskind, *op cit*, notes 13 & 18; A Terrett *The Internet: Business Strategies for Law Firms* Law Society Publishing, London, 2000.
- 21 For example, Clifford Chance's *Nextlaw*, and Linklater's *Blue Flag* (see Terrett, *op cit*, note 20).
- 22 For example, the Freeserve Desktop Lawyer available at <http://freeserve.co.uk>.