Knowledge Web: A Public Domain Expert System Delivery Environment[*]

Max Bramer
Department of Computer Science
and Software Engineering
University of Portsmouth
Portsmouth, UK
max.bramer@port.ac.uk

Abstract - *Knowledge Web is a public domain expert system delivery environment aimed at enabling a wide range of high-quality expert systems to be created and made available free of charge to users via the world wide web. Applications have a uniform 'look and feel', do not require any special software to use apart from a standard web browser and are fully integrated with other sources of information on the web. Knowledge Providers make use of a rapid application development facility which automatically converts expert system modules written in a simple scripting language into rules and questions for a production rule interpreter running on the Knowledge Web server.*

Keywords: Expert System, Delivery Environment, Production Rule Interpreter, Automatic Module Generation.

1 Introduction

The widespread availability of PCs, low cost internet access in the home and the World Wide Web has led to demand for increasingly sophisticated information (e.g. legal, financial, medical) delivered directly to the end user, not via an intermediary such as a lawyer, accountant or doctor. However the information available is not always reliable. When it is, it is generally purely static (like a textbook) and factual, often requiring substantial experience to interpret well and not tailored to the precise requirements of any specific user.

Expert Systems became popular in the 1980s and 1990s, with a large number of applications developed in many fields, often with very short development times. Despite their great potential value, the practical usefulness of many systems was severely constrained by the limited availability of computers on which to run them. For much of the time it was not even possible to assume that potential 'professional' users (such as businessmen, scientists and doctors) would have a computer to run them on in their offices, laboratories or surgeries, far less that individual users would have personal computers in their homes. The web is potentially an excellent delivery platform for expert systems able to meet the specific needs of individual users, which offers the opportunity to distribute high-quality expertise widely at low or zero cost. However, there would appear to be few expert systems available on the web. Those that are offer far less functionality than the expert systems of a quarter-century ago.

One typical web-based medical expert system asks the user around 20 yes/no/don't know questions. All users answer the same questions and the output is a probability of contracting a certain disease. There is no indication of how the final prediction was reached, how the answers given affect the results, why the questions were asked or how users might change their lifestyle to reduce the risk. The value of providing facilities to explain the purpose and meaning of questions and for the user to study the reasons underlying the system's conclusions, especially in sensitive areas such as medicine, was established as far back as the MYCIN medical consultation system in the mid-1970s. Building an expert system to run on the web, e.g. as a set of linked HTML pages, with facilities comparable to those of the expert systems of the 1980s would be a substantial task even for a single application and it is clear that the demand for such systems once established could soon become very large indeed.

Knowledge Web is a public domain expert system delivery environment aimed at tackling this problem by enabling a wide range of high-quality expert systems to be created and made available free of charge to users via the web. The design aims of the Knowledge Web project are given in the next section.

2 Knowledge Web: Design Aims

The overall aim of the Knowledge Web project is to use Information and Communication Technology to facilitate the sharing of expertise in a wide range of fields on a worldwide scale and free of charge. This is to be achieved by creating and making available a public

[*] **0-7803-7952-7/03/$17.00 © 2003 IEEE.**

domain expert system delivery environment where many people and organizations can place expert system modules for anyone to use free of charge. Ease of module development and the need to serve a very diverse community of users are two of the cornerstones of the project. The detailed design aims are as follows.

(a) Expert system modules should be made available to a large number of potential users worldwide without incurring significant expense and should be capable of being used by many people simultaneously. This strongly suggests making them available via the World Wide Web.

(b) Users should require a minimum of technical knowledge and should not need to download special software to run modules. This suggests using a standard web browser as the delivery vehicle. Browsers are already widely used worldwide and are available free of charge. The browser should not require any additional features, e.g. plugins, to be added before use.

(c) Modules should have a straightforward user interface with a uniform 'look and feel' so that user experience of the facilities available gained using one module is directly transferable to other modules. A natural way to do this is to make the output from the module appear as a succession of web pages, each one requiring some input from the user.

(d) Modules must give answers, recommendations etc. that are tailored to the individual user, i.e. the pages should not merely be pre-stored text.

(e) There has been a marked increase in the volume of regulations, codes of practice etc. in many fields in recent years. In those areas it is no longer sufficient to give the user good advice. That advice must be supported by details of regulations etc. that support it, sufficient to demonstrate that the user who follows the advice has acted in accordance with a general 'duty of care'. The provision of facilities to do this is an important part of module design.

(f) It should be possible for modules to be 'mass produced', i.e. developed rapidly and easily using special development tools. It should be possible for the modules made available by a knowledge provider to link to one another in the interest of 'knowledge reuse' and to facilitate hierarchical decomposition of complex tasks.

(g) Modules should give a rapid response. Ideally the user should not notice any difference in response from accessing a succession of standard (i.e. pre-written) web pages.

(h) Quality control should be achieved by allowing only registered 'knowledge providers' (generally, large companies and trusted individuals) to create modules.

(i) As far as possible, the user of a module should be encouraged and helped to gain understanding of and expertise in the task domain, rather than simply accept the module's decision, recommendation etc. as infallible. For this reason the module delivery environment needs to be fully integrated with the retrieval of other information from the world wide web. Where appropriate modules should support a 'mixed initiative' style of use, where the user sometimes provides information passively in response to questions and at other times seeks information more actively.

(j) Finally, issues of legal liability need to be dealt with. This is a difficult area and prominent disclaimers although probably essential are only part of the solution. Although issues of liability are a potential legal minefield which must not be underrated, the user is far less likely to be harmed or misled by an expert system module that explains its decisions and is provided in a closed delivery environment by a reputable knowledge provider than by reading through the vast quantity of frequently unreliable information on every topic that is the unrestricted world wide web.

3 Knowledge Web: System Overview

Knowledge Web comprises a number of web pages and programs written in the server-side language PHP which maintain and make use of a library of *modules* written in a specially designed production rule language. Knowledge Web has three types of user: (standard) *users* who run modules to obtain information, *knowledge providers* who create and upload modules and the *system administrator*, who maintains the module library. Facilities are provided for all three.

All Knowledge Web modules are available to anyone who wishes to use them, irrespective of which organization created them. All users need to register so that the system can distinguish between different users accessing and providing information to the same module simultaneously, but there are no restrictions on who may be a user. From the user's point of view running a module is indistinguishable from accessing a series of web pages using a standard web browser. The module asks a series of questions requiring either multiple choice answers or input of a number or string of characters. Providing this answer automatically generates a further web page with additional information and another question until the end of the consultation is reached. The modules have a uniform interface which is easy to use and can be fully integrated with additional sources of information elsewhere on the web. Further details are given in Section 4 below.

Although it would be possible for knowledge providers to create modules in the form of a possibly large number of linked web pages, this would be an onerous and error-prone task. It is an important part of the Knowledge Web design philosophy that complex modules can be created in a simple fashion. From the knowledge provider's point of view each module is written as a single plain text file in a simple scripting language (see Section 5 for details). The knowledge provider uploads the script file to the Knowledge Web website, providing additional general information during the process. One of the central components of the Knowledge Web package is a *module conversion program* which converts this script file and the other information provided during uploading into a possibly large number of rule files, question files and other files which collectively form the module. The rules and questions are written in a specially designed production rule language which is invisible to the knowledge provider as well as the user.

At each stage of a user consultation a *forward chaining production rule interpreter* (written in PHP and running on the Knowledge Web server) makes use of the rules and questions in the current modules and the series of answers supplied by the user up to that point to compute the next webpage to output to the user's browser.

As well as uploading modules knowledge providers have facilities to maintain the modules which they have previously uploaded. By default modules are uploaded in 'test' mode and can only be run by the corresponding knowledge provider. (Information about modules in test mode is not included on the menu of modules provided to users.) When satisfied with a test module the knowledge provider can re-designate it as 'live'. Knowledge providers can of course also delete their modules (either 'test' or 'live') at any time.

The system administrator has facilities to run any module, change the status of any module from live to test or vice versa or to delete any module regardless of who created it.

4 Using a Knowledge Web Module

Running a Knowledge Web module requires only a standard web browser. Having pointed his or her browser to the Knowledge Web home page and entered a valid user name and password the user is presented with a menu giving basic information about each module (name, knowledge provider, brief description etc.) such as the one shown in Figure 1.

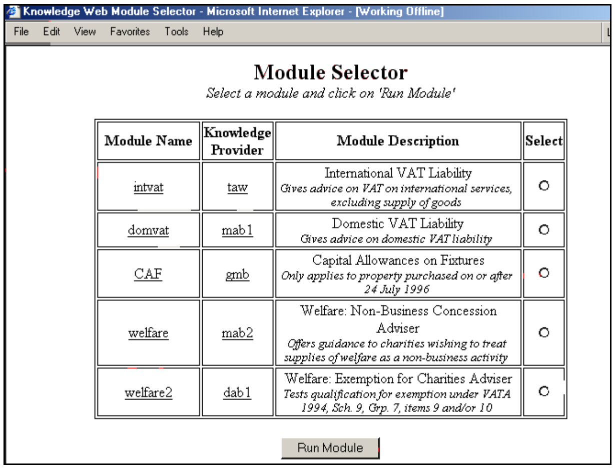
Knowledge Web Module Selector - Microsoft Internet Explorer - [Working Offline]
File Edit View Favorites Tools Help

Module Selector

Select a module and click on 'Run Module'

Module Name	Knowledge Provider	Module Description	Select
intvat	taw	International VAT Liability *Gives advice on VAT on international services, excluding supply of goods*	○
domvat	mab1	Domestic VAT Liability *Gives advice on domestic VAT liability*	○
CAF	gmb	Capital Allowances on Fixtures *Only applies to property purchased on or after 24 July 1996*	○
welfare	mab2	Welfare: Non-Business Concession Adviser *Offers guidance to charities wishing to treat supplies of welfare as a non-business activity*	○
welfare2	dab1	Welfare: Exemption for Charities Adviser *Tests qualification for exemption under VATA 1994, Sch. 9, Grp. 7, items 9 and/or 10*	○

Figure 1. Knowledge Web Module Selector

Clicking on the name of a module produces further information about that module. Clicking on the name of a Knowledge Provider gives details of the organisation or individual concerned.

Selecting a module and pressing the 'Run Module' button generates an initial frameset similar to that shown in Figure 2 with a horizontal frame giving the Knowledge Web name and logo and the name of the module, plus three vertical frames. From left to right these are the *Information Window*, the *Question Window* and the *Advice Notes Window*, respectively. The framesets for each module have a uniform style, layout, colour scheme

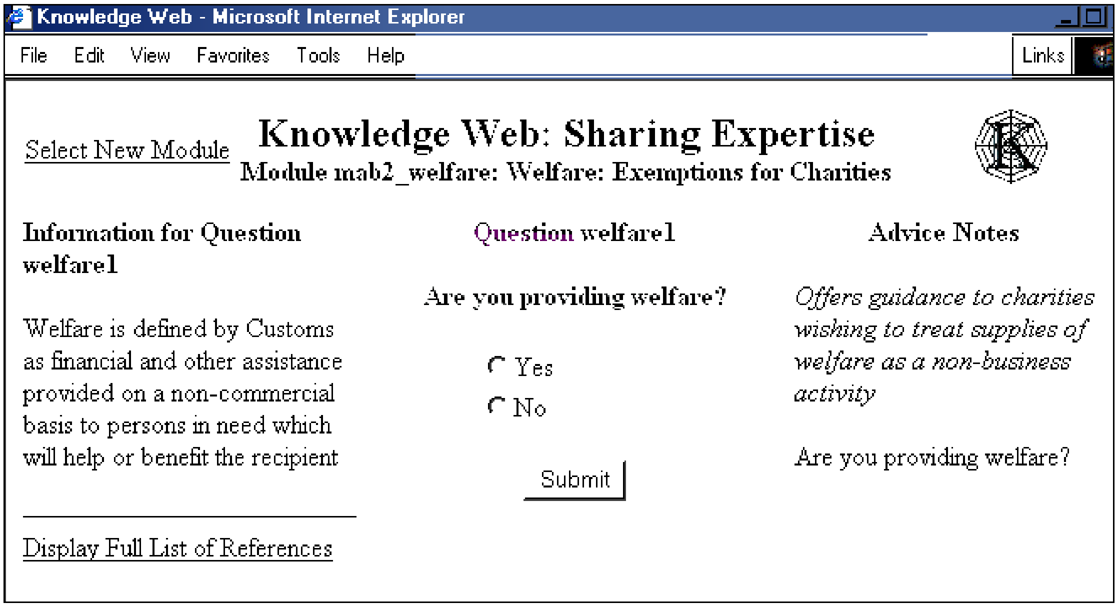
Knowledge Web - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Links
Select New Module

Knowledge Web: Sharing Expertise

Module mab2_welfare: Welfare: Exemptions for Charities

Information for Question welfare1

Welfare is defined by Customs as financial and other assistance provided on a non-commercial basis to persons in need which will help or benefit the recipient

Display Full List of References

Question welfare1

Are you providing welfare?

○ Yes

○ No

Advice Notes

Offers guidance to charities wishing to treat supplies of welfare as a non-business activity

Are you providing welfare?

Figure 2. Initial Screen for Module *welfare*

etc. (The screenshots in this paper have been converted to black and white throughout.)

The system then takes the user through a question-answer dialog, each answer resulting in a new frameset of similar appearance to before. The sequence of framesets appears to the user as a pre-stored set of web pages but is in fact generated automatically by a web based production rule interpreter, a server-side program written in PHP running on the Knowledge Web server. The processing time taken to achieve this is insignificant compared with the transmission delays that web users unavoidably experience.

The questions asked at each stage depend on previous answers, so there is no need for users to answer or even be aware of questions that are irrelevant to their own consultations. At each stage an option is available to undo the previous answer. Each question can be either multiple choice or require a number (e.g. gross income) or a line of text (e.g. user name) to be entered. If appropriate, the system can do arithmetic with the answers provided. Each question can have information associated with it (in the information window) explaining the meaning or significance of the question. This can include hot links to web pages. All links are also added to a References web page for possible future use. This facility enables very elaborate supporting information to be associated with a question (perhaps using speech, animation or video clips etc.) outside the module itself.

As a consultation proceeds information is written to the Advice Notes window. This can include the values of variables or arithmetic expressions calculated from the user input provided. The user would normally print this or paste it into a document to take away. Each piece of advice written can have associated hot links to web pages. These are also added to the References web page.

In many fields it is not sufficient for a module merely to provide the user with good advice. It is also necessary to justify it by means of references to 'authorities', such as legislation, regulations, codes of practice, EU directives etc. The information written to the Advice Window and the links to web pages placed in the Information and Advice Notes Windows and on the References web page are an important way of doing this.

Knowledge Providers are also encouraged to associate a module website with each module. A set of web pages giving answers to Frequently Asked Questions can also be associated with a module. Links to both are placed on the References web page.

Figure 3 shows the format of the screen displayed to the user at the end of a typical consultation. The Advice Notes Window gives links to two pages of further information, one on the module website, the other a FAQ page. The References page is superimposed, showing links to several web pages plus the module home page, the FAQ index page and the module search engine.

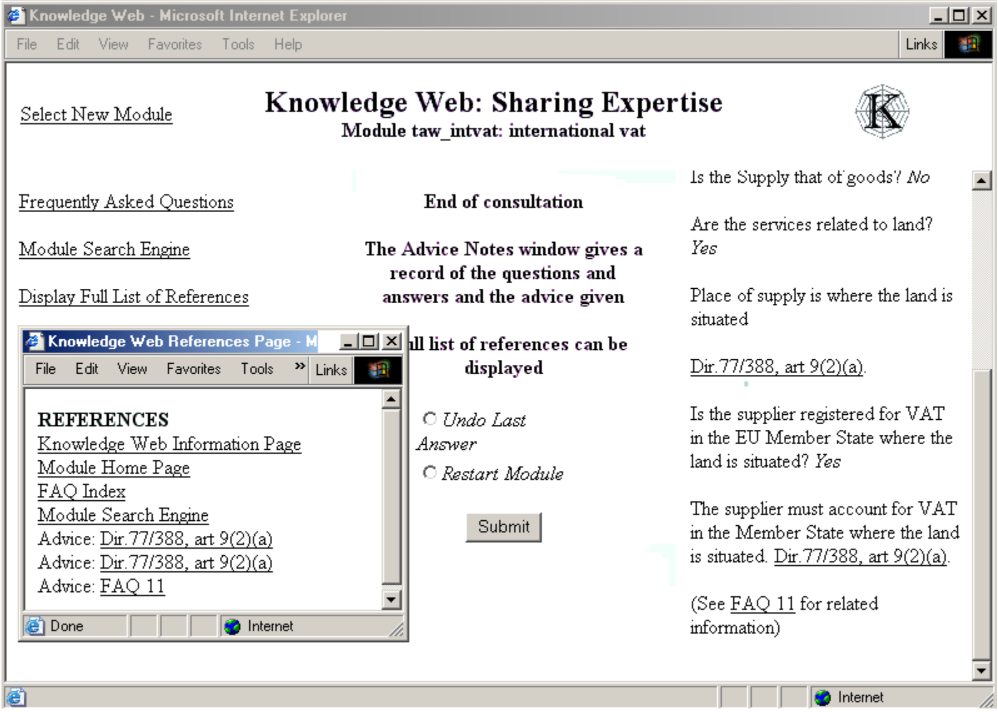
Knowledge Web - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Links
Select New Module
Knowledge Web: Sharing Expertise
Module taw_intvat: international vat
Frequently Asked Questions
Module Search Engine
Display Full List of References
Knowledge Web References Page - M
File Edit View Favorites Tools » Links
Done
Internet
○ Undo Last Answer
○ Restart Module
Submit
Internet

End of consultation

The Advice Notes window gives a record of the questions and answers and the advice given

...ll list of references can be displayed

REFERENCES
Knowledge Web Information Page
Module Home Page
FAQ Index
Module Search Engine
Advice: Dir.77/388, art 9(2)(a)
Advice: Dir.77/388, art 9(2)(a)
Advice: FAQ 11

Is the Supply that of goods? *No*

Are the services related to land? *Yes*

Place of supply is where the land is situated

Dir.77/388, art 9(2)(a).

Is the supplier registered for VAT in the EU Member State where the land is situated? *Yes*

The supplier must account for VAT in the Member State where the land is situated. Dir.77/388, art 9(2)(a).

(See FAQ 11 for related information)

Figure 3. Final Screen for a Typical Consultation

Figure 4 shows a Frequently Asked Question index page for a typical module.

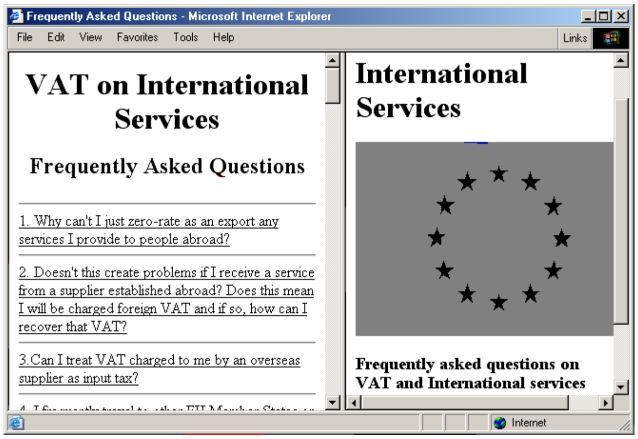
Frequently Asked Questions - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Links
Internet

VAT on International Services

Frequently Asked Questions

1. Why can't I just zero-rate as an export any services I provide to people abroad?

2. Doesn't this create problems if I receive a service from a supplier established abroad? Does this mean I will be charged foreign VAT and if so, how can I recover that VAT?

3.Can I treat VAT charged to me by an overseas supplier as input tax?

International Services

Frequently asked questions on VAT and International services

Figure 4. FAQ Page for Module INTVAT

The Knowledge Provider can also make a module search engine available to the user which will perform a keyword search beginning at the module home page and following all links through until some preset search limit is reached.

It is important for the modules to have the facility to be fully integrated with 'normal' web pages, but the need to use this facility depends on the task area concerned. At one extreme, it is possible to write modules with no web links, no home page etc. At the other, the module might be just one or two questions that act as a simple front end to an extensive website. In general the system will be 'mixed initiative' with the user able to ignore the questions and go into web browsing or FAQ reading mode whenever he/she wishes. One of the less obvious but still important design aims is that using the system will also help to educate the user about the task domain.

5 Knowledge Web Script Files

The script language used by Knowledge Web is a much augmented form of the scripting language originally provided for the SPSS Inc. flowcharting package allClear 3.5.

Figure 5 shows a very basic script file. The alternative answers to each question are indented from it and enclosed in parentheses. (Thus the indentation shows the linkage between questions and answers.) Lines such as 'Concession applies' are called *advice*. Any number of advice lines can be placed after an answer either before or instead of a further question. The line -lab1 is a *label*, which can be jumped to from four places in the script (the >lab1 lines).

Even a very basic script such as this will be converted automatically by the Knowledge Web conversion program into a module of the kind shown in Section 4. At each stage the questions and alternative answers will automatically be placed in the Question Window and the advice lines will be placed in the Advice Notes Window.

```
Are you providing welfare?
    (Yes)
    Is service to an individual?
        (Yes)
        Is service available to all distressed people?
            (Yes)
            Is the service subsidised > 15% below cost?
                (Yes)
                Is it your policy to subsidise the service?
                    (Yes)
                    Concession applies
                    (No)
                    -lab1
                    Concession does not apply
                (No)
                > lab1
            (No)
            > lab1
        (No)
        > lab1
    (No)
    > lab1
```

Figure 5. A Basic Script File

To enable a much richer variety of information to be made available to the user the basic form of script can be extended in many ways. Tags similar to those used by HTML can be used to associate background information with any question and to link any number of web references to any question or any advice text. In addition, modules can easily be linked to other modules, enabling top-down decomposition of a complex task into a number of smaller ones. Users can be asked to provide input (either numerical or text) as well as to choose from a set of alternative answers. Values can be assigned to variables and the value of an arithmetic expression can be evaluated and either assigned to a variable or output. Further details of these and other facilities are given in the following sections.

Experience has shown that creating these script files is straightforward and takes far less time than writing a set of connected web pages to cover all possible combinations of user responses. The conversion program takes only a few seconds to convert even the most complex script into the possibly hundreds of files that comprise a module.

5.1 Associating Information with Questions

Background information can be associated with any question simply by appending the *information tag* *INF* followed by the information text to the end of the question, e.g.

```
Are the services related to land?*INF*YES
answers include:-The Grant, assignment or
surrender of any interest in or right over land
```

When the question is asked during a Knowledge Web consultation the text following the information tag will automatically be placed in the Information Window to guide the user in answering the question.

5.2 Associating Web References with Questions

A web reference can be associated with any question by placing the *reference tag* *REF* followed by a *reference string*, followed by two asterisks and a URL at the end of the question, e.g.

```
Is the supplier registered for VAT in the EU
Member State where the service is
performed?*REF*Dir 77/388 Art 9(2)(c)**6dir-
9.htm
```

In this example a hyperlink with the text *Dir 77/388 Art 9(2)(c)* will be placed on the References page when the question is asked during a consultation. If the user later clicks on that menu entry the corresponding web page will be displayed.

Both background information and an unlimited number of web references can be associated with the same question, e.g.

```
Are the services related to land?*INF*YES answers
include:-The Grant, assignment or surrender of any
interest in or right over land*REF*Directive 98/45
section 3**sect3.htm*REF*Guidance Note 98/47
Section 6**note98_47.htm
```

Note that URLs can be entered in an abbreviated form as above, e.g. *sect3.htm*. Any URL not beginning with http:// will automatically be prefixed by the module 'web stem' (see Section 5.7 below).

5.3 Associating Web References with Advice Text

As in the case of questions, web references can be associated with any 'advice text' in a flowchart by placing the reference tag *REF* followed by a reference string, followed by two asterisks and a web URL at the end of the advice text, e.g.

```
VAT is chargeable*REF*Dir.77/388, art. 9(2)(a)
**6dir-9.htm
```

The reference string *Dir.77/388, art. 9(2)(a)* will be placed on the References webpage when the corresponding piece of advice is displayed in the Advice Notes Window during a consultation. If the user then clicks on that menu entry the corresponding web page will be displayed.

It is possible to associate an unlimited number of web references with any advice text, e.g.

```
The supplier must account for VAT in the Member
State where the land is situated. Dir.77/388, art
9(2)(a)*REF*Dir.77/388, art 9(2)(a)**6dir-
9.htm*REF*FAQ 11**faq/Q11.htm
```

As in the previous section URLs can be entered in an abbreviated form and will be automatically prefixed by the module webstem.

5.4 Automatic Generation of Hot Links

If there is any text in either the Information Window or the Advice Notes Window which exactly matches a web reference currently on the References webpage the conversion program will automatically generate a hot link in the corresponding window.

The relevant text will automatically be displayed in colour and underlined in the usual way, e.g.

> See FAQ 10 for information about cultural and artistic services
>
> See FAQ 26 and FAQ 27 for information about ancillary transport services

Clicking on any of the underlined text will then cause the corresponding web page to be displayed.

Where a web reference is associated with a question but no background information text is given, the conversion program will automatically generate an appropriate information text to ensure that there is a hot link from the Information Window, e.g.

> Additional information about Section 145(1) ICTA 1988 is available either from the References page or by double clicking on the underlined text

5.5 Linking Modules

At any point where a piece of advice would be appropriate a script can include a call to another module created by the same Knowledge Provider, e.g.

```
*ATTACH* carbens
```

If this point is reached during a Knowledge Web consultation execution of the module will cease and a new module will automatically be started.

This is a very powerful facility which enables large applications to be constructed by a process of hierarchical decomposition.

Developers of a module can choose names for variables and labels without knowing which choices have been made by other developers. Knowledge Web distinguishes between occurrences of the same name in different modules.

5.6 Using Variables

Variables can be given values in three ways in a Knowledge Web script.

(a) A standard (multiple choice) question can be augmented using the *variable tag* *VAR* to assign the value chosen by the user to a specified variable, e.g.

```
What is the VAT rate?*VAR*vrate
```

(b) An *input statement* can be used anywhere in a script file where a piece of advice would be valid to ask the user to enter a value, e.g.

```
INPUT y * What is the value of y?
```

Once a value has been entered it is automatically assigned to the specified variable.

(c) A *let statement* can be used anywhere in a script file where a piece of advice would be valid to give a variable a value. This can be either a constant or the value of an arithmetic expression, e.g.

```
LET z = (x+y)/3
LET w = z-y-4.2+2*x
```

Both variables and arithmetic expressions can be output simply by being included in a piece of advice. They are distinguished from standard text by being enclosed in square brackets, e.g.

```
Your profit of [x] pounds will incur a tax liability of
[(x-allowance)*rate] pounds
```

5.7 Information Provided During Uploading

Although not part of the script file itself the information entered by the Knowledge Provider when uploading the file complements that in the script file in several ways.

(a) The Knowledge Provider must enter a module name and also a short phrase describing the module. Both are entered in the module selection table displayed to users. Optionally a longer *module description* (up to 240 characters) can also be entered. If it is, it provides the additional information about the module available to the user from the module selection table mentioned in Section 4 and is also automatically placed at the head of the Knowledge Web Advice Notes Window as an introductory paragraph about the module.

(b) Where appropriate the Knowledge Provider can specify the URL of a home web page for the module and also a web page of Frequently Asked Questions (FAQs). If specified these are automatically added to the References page whenever the module is used.

(c) The Knowledge Provider can also specify whether a module search engine should be made available to the user. This will normally only be worthwhile if there is a reasonably extensive module website or set of FAQ pages. The search engine enables the user to perform a keyword search for information, but for quality control reasons this is restricted to the module home page and those pages linked directly or indirectly to it.

(d) Finally, to reduce the need for possibly error-prone typing of long URLs the Knowledge Provider can specify a web stem to be automatically prefixed to any URLs associated with either questions or advice text that do not begin with http://.

6 Conclusions

Knowledge Web is currently undergoing trials. A range of large modules has successfully been created using automatic conversion from script files for purposes of evaluation. The largest suite of modules links together approximately 20 modules written independently.

Feedback from developmental testing has been favorable and discussions are currently taking place with potential knowledge providers for the live system.