Chapter 23

MAKING DECISIONS ABOUT LEGAL RESPONSES TO CYBER ATTACKS

L. Peng, T. Wingfield, D. Wijesekera, E. Frye, R. Jackson and J. Michael

Abstract Cyber intrusions may be characterized in one or more of three legal regimes: law enforcement, intelligence collection and military operations. Furthermore, most intrusions occur across a number of jurisdictional boundaries, building complex conflict-of-laws questions into such attacks. Applying a one-size-fits-all response, such as always terminating all interaction with the intruder or always responding in kind, can be an ineffective or even worse, illegal, response. In order to assist investigators and legal experts addressing the legal aspects of cyber incidents, we have developed a decision support tool that takes them through a series of questions that are akin to those posed by an attorney to a client seeking legal guidance. Our tool may be used by builders and users. Builders use the tool to construct trees of legal arguments applied to the incidents at hand with the documentation useful for building legal briefs. Users interact with the tool by answering a series of questions to obtain viable legal arguments with supporting documents.

Keywords: Cyber attacks, legal issues, decision support system, incident response

1. Introduction

When a malefactor intrudes into a computer system, the owner of that system – whether a private individual protecting personal property, a corporation securing its assets, or a government defending its interests – needs to know something about the malefactor in order to develop a lawful and effective response to the intrusion. Cyber intrusions may be characterized in one or more of three legal regimes: law enforcement, intelligence collection and military operations. Furthermore, intrusions can occur across a number of jurisdictional boundaries, building complex conflict-of-laws questions into such attacks. Applying a one-size-fits-all response, such as always terminating all interaction with the intruder or always responding in kind, can be an ineffective or even worse, illegal, response. For instance, terminating interaction with an intruder could prevent the seizure of evidence for criminal prosecution, collection of information for counterintelligence purposes, or counter-targeting for a military response [9]. By responding in kind, the defender may violate domestic or international law, or, in the case of a government actor, inadvertently escalate to the level of a use of force or even an armed attack. Furthermore, an inappropriately calibrated response may contravene the customary rules of war accepted as authoritative law by the United States – distinction, necessity, proportionality and chivalry.

The general problem we address in this paper is that of providing defenders with sufficient information to make informed decisions when formulating responses to intruders. Specifically, we describe a tool that serves as an automated aid for determining the legal regime under which a cyber intrusion can be categorized, with all documentation supporting the building of a brief. Our tool is built on the premise that owners and their agents of affected computing resources want to defend their computer systems without violating domestic or international law.

Both the frequency and intensity of attacks in cyberspace can be high, affording little time for research and thoughtful consideration before the cyber intrusion (whether a theft or an attack) is over. Similarly, what may initially appear to be a minor intrusion or misuse of a computer system may ultimately result in damage to or destruction of property, or even human injury or loss of life. In either case, the defender must be prepared to respond to such attacks with operational plans and mechanisms for real-time information collection already in place. In other words, the defender needs to tighten his Observe-Orient-Decide-Act (OODA) loop in order to gain a competitive advantage over the intruder [7].

Legal preparation is an essential element in this equation. Against opponents who disregard any laws which are not immediately and effectively punitive, the default response of inadequately counseled operators is to forego otherwise lawful and effective defensive strategies. In other words, the vast legal gray area that exists today operates in favor of the intruder – a form of asymmetry between the attacker and the defender. A clearer and more timely picture of the operational legalities of the situation would provide the defender with more, rather than fewer, options. At this stage in our research, several caveats are in order. First, the present tool is illustrative of the concept, and is not intended to be employed operationally at this point. The questions and answers have an artificially academic clarity, which derives from top-down reasoning of broad questions to narrow circumstances. Second, the decision-tree format no longer defines the state-of-the art in expert systems, but it does: (a) present the core concepts clearly, (b) provide a framework that clarifies the transparent assembly of resources supporting legal analyses, and (c) lay a foundation for more elaborate logical structures (such as totality-of-the-circumstances analysis) for future operational employment. Third, the inevitable anomalies which will arise in its development (i.e., requiring an early legal determination of whether or not the intruder is a US person will almost certainly conflict with the operational reality of discovering key facts late in the game) serve to highlight conflicts and lacunae in the law. The degree to which the most operationally useful flow of legal questions fails to meet real world requirements is the degree to which the law or technology must change. Fourth, this tool will be developed in alignment with international law, but numerous questions (especially in the law enforcement and intelligence collection realms) will never rise to the level of state vs. state legal determinations. Where national and international law appear to conflict, that tension will be made explicit and thus clarified for resolution.

The remainder of this paper is organized as follows. Section 2 describes the details of our application requirements. Section 3 describes the software design of our toolkit. Section 4 explains the functionality of the tool through an example. Section 5 describes related work and Section 6 concludes the paper.

2. Legal Requirements

As stated, our objective is to enforce legal responses to cyber incidents. In doing so, we are guided by the legal advice given to litigants who claim that they have a legitimate case for recourse. When a litigant discusses his or her situation with an attorney or investigator, the latter asks a series of questions to determine the applicable legal regime and to map out a course of action. Our larger objective is to make this a primary global requirement for responding to incidents in a timely manner. To reason about response alternatives, we first need a model of the domestic and international law governing cyber intrusions, one for computers to execute without the human in the loop and at high speed, and another for human decision making at considerably lower speed. Our proposal for this model is a customizable decision tree of legally relevant questions, modeling those that would be asked by an investigator from a prospective complainant. While the computer's decision capability encoded in the form of a tree can be hardwired for independent execution of clearly discernable, objectively verifiable criteria, the decision tree to be manually traversed will have pre-selected sources available to assist the attorney in deciding each of the gray area judgments requiring human reflection and creativity. It is necessary to assemble a comprehensive selection of sources to append these to each decision point, but it will be vital, for speed and clarity, to include no more than is required to answer the question at hand. These sources may be grouped as constitutional, legislative (statutes), executive (regulations), judiciary (cases) and international. These five categories must be further subdivided into primary (the case or statute itself), and secondary (analytic and synthetic commentary, such as law review articles, or briefs on file). These ten categories are sufficient to contain any legal source needed to address any given question. Furthermore, each source would have to be presented at four levels of abstraction, for the proper balance of speed and depth:

- Citation: A legal footnote.
- Précis: A sentence or paragraph paraphrasing what the source has to say about the question at hand.
- Excerpt: Direct quotes from the source which are on point
- Document: The complete law review article, statute or case.

This general information would be distilled into a specific research question in two media: an audit trail, providing a record of each question asked and each answer chosen, and a brief builder, which would augment the audit trail with those portions of the sources selected by the reviewing attorney to support his answer to the question. This would, in effect, be the first draft of a legal brief supporting the selected course of action.

The decision tree, and its supporting sources, may be constructed using an open source methodology, allowing law students, practitioners, and scholars scattered across the world to collaborate on its construction and refinement. With the process architecture (described below) in place, the trees will be available to selected legal academics for analysis and improvement. Designing such a legal analysis tool for a comprehensive tracking system will be of great benefit to the cyber-legal community, because it will require the analysis and distillation of the entire field into the simplest possible framework for implementation.

This system will take the form of a set of predefined sequential questions when an actor's behavior indicates he or she may be intruding into, misusing or attacking a computer system. To simplify the logic employed, in the prototype, each question has only *yes* and *no* answers. A deferent question will follow each *yes* or *no* answer to continue the

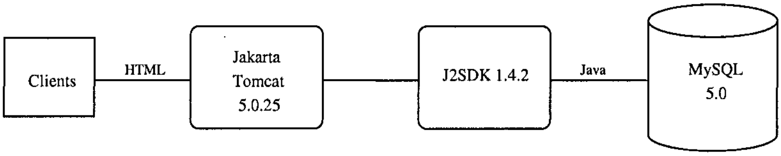

Figure 1. System architecture.

analysis. Then attorneys and their clients would follow a complete logical path to reach a transparently reasoned legal conclusion. A third option, *don't know* allows the user to view the legal resources necessary to proceed forward with a *yes* or *no* answer. As mentioned earlier, these resources are arrayed in ten categories (constitutional, legislative, executive, judicial and international, each at a primary and secondary level), and each source may be accessed at any one of four levels of abstraction (citation, précis, quotation and full source).

The system will operate on two levels: for users following previously constructed analyses, and for builders, assembling and testing the analyses to be provided to the operational community. Users are attorneys responsible for providing operational legal advice to law enforcement, intelligence community, or military officials. These users follow a decision tree, answering a sequence of questions carefully crafted to identify and record the legally operative facts of the incident. This decision support tool will produce a logical legal analysis, supported by the legal resources selected by the user. Builders are academic or practicing attorneys and some computer network technicians, adding and subtracting branches from the decision tree and the resources available at each decision point. They create and maintain the substance of the decision support tool.

3. System Design

The prototype is designed to be an open-source, web-enabled decision support tool that provides legal reasoning web services. Multiple clients may access the web server (the system) via web browsers, such as Internet Explorer or Netscape. The communication language between clients and web server is HTML exposed within Java Server Pages (JSP). A Java engine, Java 2 Software Development Kit (J2SDK), is used to compile JSP pages to Java class files that send a HTML stream to web clients and communicate with a mySQL database through JConnector technology. Figure 1 shows the system architecture. Compared to client-server applications, among others, this multi-tier design has the following advantages:

- Clients may remotely and concurrently access the system, sharing the same knowledge base.

- The architecture is extensible, because it is built using Java 2 Enterprise Edition (J2EE) service framework, with quick deployment times and minimal maintenance efforts in mind.

- The system can be extended to use RDF, OWL, RuleML or JESS as needed.

- The system is easily manageable, because some clients are allowed to change the knowledge base while other clients can only access built-in scenarios.

Each client (actor in software engineering) is a builder or user with his/her own separate applications that share one database and file system. User functionality includes answering questions, getting a decision, viewing (audit trail, tree map, legal brief), searching pertinent legal documents, and displaying legal documents. Builder functionality includes adding and deleting trees/decisions/questions, linking decisions/resources and loading resources.

3.1 Detailed User Requirements

- The system should collect legally relevant facts

- The system should follow the decision tree, answering *yes, no* or *don't know* to each question in sequence. A user should be able to go back to prior questions and change answers to evaluate the consequences of alternative answers. An *yes* or *no* answer proceeds to the next (pre-determined) question. While many possible paths are available, any given sequence of *yes* or *no* answers should yield only one result.

- For a *don't know* answer, the system should present legal resources to assist in making a *yes* or *no* decision. These resources will vary in number and length depending on the question at hand, but are grouped by category (constitutional, legislative, executive, judicial and international), each with a primary and secondary set of materials, and subcategories such as country and language. Each resource is accessible to four levels of detail: citation, précis, excerpt and source.

- The audit trail function should display the history of navigation with consulted sources in the citation format along with answers provided by the user.

- The brief builder function should do the same, and include all user-selected portions of consulted sources.
- After a sufficient number of questions have been answered, the system should provide a decision with supporting documents. A user should be able to search databases uploaded into the tool under the ten categories.
- The system should display searched resources at four levels of detail (citation, pécis, excerpt and source).

3.2 Detailed Builder Requirements

- Builders should be able to login and navigate any of the tool's web pages, including those of users. For security reasons, builders inactive for thirty minutes are logged out.
- After initiating a new decision tree or selecting an existing one, a builder should be able to add any answer *(yes/no/don't know)* and link either of these to another question.
- Builders should be able to upload relevant documents and categorize them in support of the *don't know* option.
- Builders should be able to separate each resource into its appropriate levels of abstraction (citation, pécis, excerpt and source).
- Builders should be able to delete a question, a decision or an entire tree.

4. System Functionality

This section describes the functionality of the tool by constructing an example decision tree to determine the answer to the legal question:*Are we at war?* (see Figure 2). The builder can access the system to build a decision tree via a web browser after a correct login as shown in Figure 3.

As a first step, the builder creates a new tree named: *Are we at war?* Then, the builder adds three possible decisions to this new tree. Next, the builder inserts multiple questions and links the right follow-up questions or decisions with them. The builder needs to specify the parent question in the tree that to which the new question is to be linked; that is, the builder should design the system so that a *yes* or *no* answer to a previous question is linked to a new question posed to a user as shown in Figure 4. Because decision trees can be complex, the system is designed to offer the builder flexibility. For example, the builder can input the system's decision and questions without having to enter the links when

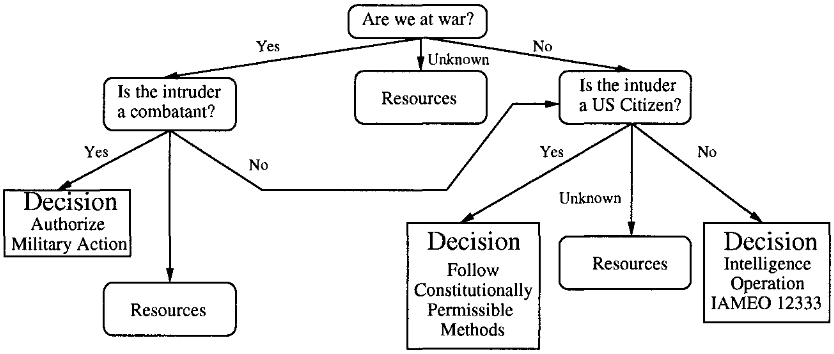
Are we at war?
Yes
Unknown
No
Is the intruder a combatant?
Resources
Is the intuder a US Citizen?
Yes
No
Yes
No
Decision
Authorize Military Action
Unknown
Resources
Decision
Follow Constitutionally Permissible Methods
Resources
Decision
Intelligence Operation IAMEO 12333

Figure 2. Decision tree.

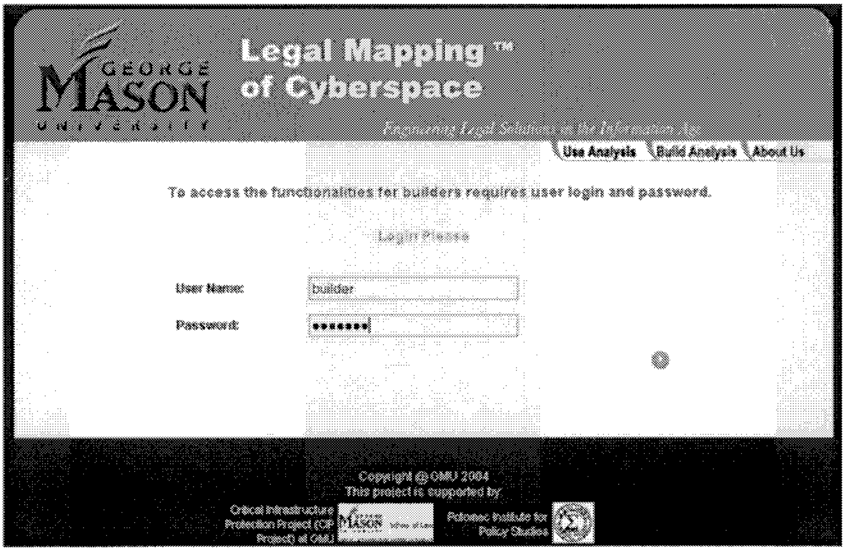
GEORGE MASON UNIVERSITY
Legal Mapping ™ of Cyberspace
Engineering Legal Solutions in the Information Age
Use Analysis
Build Analysis
About Us

To access the functionalities for builders requires user login and password.

Login Please

User Name: builder

Password: •••••••

Copyright @ GMU 2004
This project is supported by
Critical Infrastructure Protection Project (CIP Project) at GMU
Potomac Institute for Policy Studies

Figure 3. Login page of the tool.

specifying them. After that, the builder can use menu options to link the decisions and questions. The complete decision tree can be constructed in multiple ways as shown in Figure 5.

5. Related Work

Although there have been numerous academic attempts to elicit a logical structure from legal decision-making processes, none is in widespread use with practicing attorneys. The proprietary databases of Westlaw and Lexis-Nexis, searchable by document category and Boolean keyword strings, are the most frequently consulted by attorneys. Both have an

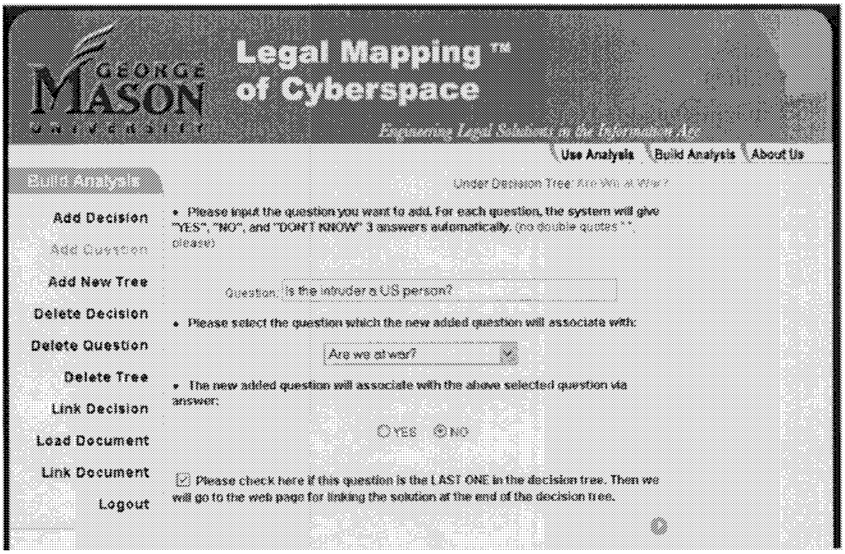

Under Decision Tree: Are We at War?

- Please input the question you want to add. For each question, the system will give "YES", "NO", and "DON'T KNOW" 3 answers automatically. (no double quotes " ", please)

Question: Is the intuder a US person?

- Please select the question which the new added question will associate with:

Are we at war?

- The new added question will associate with the above selected question via answer:

○ YES ● NO

☑ Please check here if this question is the LAST ONE in the decision tree. Then we will go to the web page for linking the solution at the end of the decision tree.

Figure 4. Linking questions.

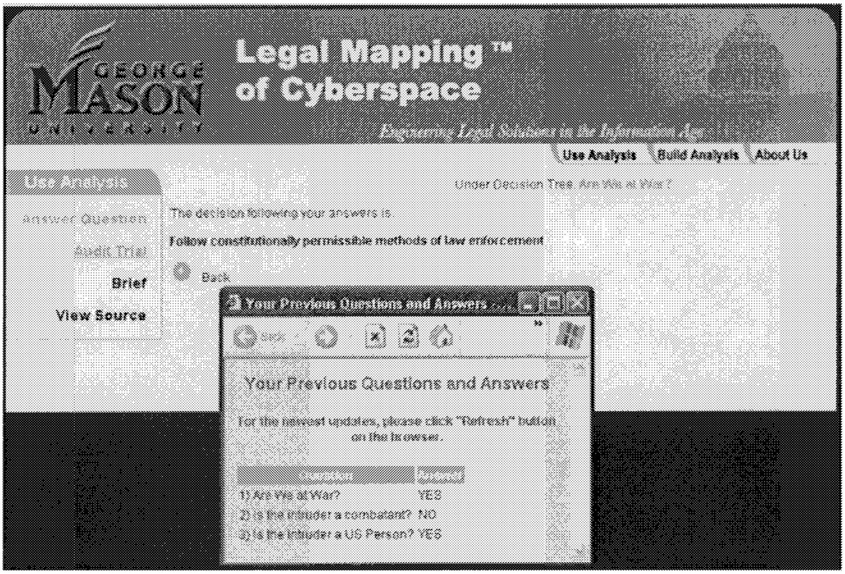

Under Decision Tree: Are We at War?

The decision following your answers is:

Follow constitutionally permissible methods of law enforcement

Your Previous Questions and Answers

For the newest updates, please click "Refresh" button on the browser.

Question	Answer
1) Are We at War?	YES
2) Is the intuder a combatant?	NO
3) Is the intuder a US Person?	YES

Figure 5. Decision rendered by a completed tree.

impressive number of cases, briefs, law review articles, and related documents [4, 5, 8], but neither is intended to provide direct assistance with the formulation or execution of a legal analysis. Furthermore, there are numerous *free* sites on the Internet – mostly maintained by universities – that have large searchable databases. Like their commercial analogs, they provide quick and reliable access to the documents themselves, but are not designed to assist in legal analyses *per se*. The University of Minnesota's Human Rights Library is an excellent example of such a system; it is the source of U.N. Charter text provided in one of our examples [10].

Capturing legal knowledge and enabling legal discourse is technically challenging and a continuous effort because laws and their interpretations change over time. Several legal reasoning support tools, e.g., [11, 12, 13], are used primarily by law students to hone their analytical skills. Others are geared for methodology or ontology research [3, 6, 13]. Only a few of these are complete web-based tools used for general legal reasoning [1, 14, 15], and therefore are not specific to one area of law. Digud [2] and Zeleznikow [16] have developed web-based tools for reasoning about divorce laws and enhancing access to legal process. In contrast, our tool may be used to train law students and cyberspace technicians, as well as to provide legal support for responding to cyber intrusions. Being both web-based and open source increases its usability, extensibility, maintainability and potential for incremental enhancements.

6. Conclusions

Due to the need to keep responses to cyber attacks legal, responders need to be aware of the legal support available within a given legal framework. To address this need, we have developed a decision-tree-based tool that takes potential investigators and attorneys through a series of questions to help build legal briefs against perpetrators. In order to do so, the decisions have to be constructed by attorneys who are well versed in this area of law: they construct trees of sequentially-ordered questions that guide users through to an actionable recommendations for response (i.e., answers presented at terminal leaves in a tree). In addition, our toolkit stores relevant information within legal categories (e.g., constitutional, legal, international) at four levels of detail (footnote, précis, excerpt, entire document) necessary to build legal briefs.

To improve the usefulness of our tool and serve a diverse community of international users, we are in the process of populating it with legal documents related to critical infrastructure protection from different countries. Because different legal systems use different ontologies, we are designing an interoperable abstract super-ontology that specializes to different legal systems. The super-ontology facilitates semantic searching across legal ontologies in addition to introducing a degree of transparency in that the user does not need to be familiar with a foreign legal system. For the aforementioned purpose, we follow a two-pronged approach. For the short term, we analyze international legal documents and categorize their legal discourse ontologies. For the long term, we are developing an abstract ontology that can be adapted to the specific doctrines that are used in various areas of the law.

References

[1] K. Curran and L. Higgins, A legal information retrieval system, *Journal of Information, Law and Technology*, vol. 3, 2000.

[2] S. Duguid, L. Edwards and J. Kingston, A web-based decision support system for divorce lawyers, *Journal of Law, Computers and Technology*, vol. 15, pp. 265-280, 2001.

[3] A. Gangemi, A. Prisco, M. Sagri, G. Steve and D. Tiscornia, Some ontological tools to support legal regulatory compliance, *Proceedings of the Workshop on Regulatory Ontologies and the Modeling of Complaint Regulations*, *Lecture Notes in Computer Science (Vol. 2889)*, Springer-Verlag, Berlin Heidelberg, Germany, pp. 607-620, 2003.

[4] C. Hafner, Legal reasoning models, in *International Encyclopedia of the Social and Behavioral Sciences*, Elsevier, Amsterdam, The Netherlands, 2001.

[5] C. Hafner and D. Berman, The role of context in case-based legal reasoning: Teleological, temporal and procedural, *Artificial Intelligence and Law*, vol. 10, pp. 19-64, 2002.

[6] M. Hall, A. Stranieri and J. Zeleznikow, A strategy for evaluating web-based decision support systems, *Proceedings of the Sixth East-European Conference on Advances in Data Information Systems*, 2002.

[7] P. Huygen, Use of Bayesian belief networks in legal reasoning, *Proceedings of the Seventeenth British and Irish Legal Education Technology Association Conference*, 2002.

[8] E. Katsh and J. Rifkin, *Online Dispute Resolution: Resolving Conflicts in Cyberspace*, Jossey-Bass, San Francisco, California, 2001.

[9] J. Michael, On the response policy of software decoys: Conducting software-based deception in the cyber battlespace, *Proceedings of the Twenty-Sixth Computer Software and Applications Conference*, pp. 957-962, 2002.

[10] J. Michael and T. Wingfield, Lawful cyber decoy policy, in *Security and Privacy in the Age of Uncertainty*, D. Gritzalis, *et al.*, (Eds.), Kluwer, Boston, Massachusetts, pp. 483-488, 2003.

[11] A. Muntjewerff, Automated training of legal reasoning, *Proceedings of the Ninth British and Irish Legal Education Technology Association Conference*, pp. 51-58, 1994.

[12] A. Muntjewerff, A. Jordaans, R. Huekstra and R. Leenes, Case analysis and storage environment (case), *JURIX*, 2002.

[13] V. Randall, Online academic assistance for law students (academic.udayton.edu/legaled/online/).

[14] A. Stranieri, J. Yearwood and J. Zeleznikow, Tools for placing legal decision support systems on the World-Wide Web, *Proceedings of the Eighth International Conference on Artificial Intelligence and Law*, pp. 206-214, 2001.

[15] A. Stranieri and J. Zeleznikow, Tools for intelligent decision support system development in the legal domain, *Proceedings of the Twelfth IEEE International Conference on Tools with Artificial Intelligence*, pp. 186-189, 2000.

[16] J. Zeleznikow, Using web-based legal decision support systems to improve access to justice, *Journal of Information and Communication Technology Law*, 2002.