

SHYSTER-MYCIN: A Hybrid Legal Expert System*

Thomas A. O’Callaghan
Dept of Computer Science
Australian National University
Canberra, Australia
tom@tom-ocallaghan.com

James Pople
Dept of Computer Science
Australian National University
Canberra, Australia
james@popple.net

Eric McCreath
Dept of Computer Science
Australian National University
Canberra, Australia
eric.mccreath@anu.edu.au

ABSTRACT

SHYSTER-MYCIN combines a case-based legal expert system (SHYSTER) with a rule-based expert system (MYCIN) to form a hybrid legal expert system. MYCIN’s reporting has been improved for use with SHYSTER-MYCIN to provide more useful information about the system’s conclusions.

SHYSTER-MYCIN’s output was tested against that of a group of lawyers, not expert in the test domain (Australian copyright law). This allowed the system’s reasoning, rather than its depth of knowledge, to be tested. Testing indicates that SHYSTER-MYCIN’s approach to the law—using a rule-based system to reason with legislation and a case-based system to reason with cases—is appropriate.

1. INTRODUCTION

SHYSTER-MYCIN is a legal expert system, combining two other expert systems: SHYSTER (a legal expert system) and MYCIN (a medical expert system). SHYSTER [5] is a case-based reasoner, but was designed to form part of a hybrid system. This hybridisation was achieved by the addition of the rule-based reasoner MYCIN [2].¹

The SHYSTER part of SHYSTER-MYCIN was not modified; the MYCIN part was modified, as explained below. To test SHYSTER-MYCIN, the SHYSTER part used a case law specification of the meaning of the term *authorization* in the Australian *Copyright Act 1968*; the MYCIN part was populated with rules derived from sections of that Act.

This approach of separating knowledge of cases and legislation into a case-based and a rule-based reasoner, respectively, is jurisprudentially justified.²

*Further information on the research described in this paper is available in [3] and [4], and at <http://cs.anu.edu.au/software/shyster/tom/>.

¹The version of MYCIN used for SHYSTER-MYCIN is available at <http://www.norvig.com/paip.html>. For an account of the original MYCIN system, see [1].

²For a full discussion of the jurisprudential justification, see [3] and [4].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICAIL ’03, 24-28 June, 2003, Edinburgh, Scotland, UK
Copyright 2003 ACM 1-58113-747-8. \$5.00.

2. ALTERATIONS TO MYCIN

In order to make MYCIN part of a hybrid legal expert system, some changes were required. MYCIN was stripped of its medical knowledge, and a new rule base was created: 273 rules were written to represent those provisions of the Act whose application might require a determination as to whether copyright infringement has been “authorized” for the purposes of the Act. (The SHYSTER part was used to decide whether infringement had been authorised.) 56 parameters were used to represent the facts to which these rules were applied.

The MYCIN reasoner was altered so that, when each question was asked, information about the rule then under consideration was recorded in a file. This record—called a “stream of consciousness”—detailed why the MYCIN part was asking each question, and provided information as to how the system arrived at its conclusions. This record was useful in debugging the rule base.

MYCIN’s reporting was also improved for SHYSTER-MYCIN. When the reasoner concludes a rule, information is written to a file explaining the facts that are known, the rule that was applied to those facts, and the conclusion that was made. This information was marked-up with \LaTeX tags so that it could be easily combined with the output from SHYSTER, which is also in \LaTeX format.

3. TESTING

A novel method was adopted to test SHYSTER-MYCIN. Five questions in the area of copyright law were given to the system, and to a test group of lawyers, and the reports of the system and of the test group were compared. The test group consisted of a law graduate, a practising solicitor with five years’ experience, and another with 30 years’ experience.

Importantly, none of the test group was expert in copyright law. However, each member of the test group had been trained in how to reason with the law, and all had varying levels of experience in applying and developing their methods of reasoning.

The material on which the test group were to base their decisions was the same material encoded in the knowledge bases of SHYSTER-MYCIN. That is, the group were provided with relevant provisions of the Copyright Act, and relevant case summaries and “fact-vectors” from SHYSTER’s case-base. This created a “playing field”—for SHYSTER-MYCIN and the test group—that was as close to level as possible.

This approach to testing differs from previous evaluations of legal expert systems in that SHYSTER-MYCIN's reasoning, and not its depth of knowledge, is being evaluated. When an expert system is made to compete with a human expert, the human expert can draw upon years of experience; the expert system can only draw upon what has been fed into it.

With the system and the test group working from the same level of "knowledge" about copyright law, the quality of their respective reasoning about the test questions can be compared. This is achieved by use of the "correctness" measure and, to a lesser extent, the validity measure (discussed below).

4. RESULTS

SHYSTER-MYCIN's output was evaluated against three criteria: validity, conciseness and "correctness". Evaluation was by comparison with the output from the test group. A report from SHYSTER-MYCIN was considered valid if it relied upon the same sections of the Act as did a majority of the test group. Conciseness was measured by the number of conclusions stated, rather than the overall length of those conclusions. A report was considered "correct" if the result produced by SHYSTER-MYCIN was the same as that predicted by a majority of the test group.

Different methods of reporting were experimented with during the development of SHYSTER-MYCIN and the different methods used by different versions of the system were evaluated and compared.

4.1 Validity

The previous version of SHYSTER-MYCIN ("SM-v2") reported on every section upon which it based a conclusion, which was invariably every section that it knew about. The current version ("SM-v3") takes a selective approach to its reporting: conclusions that are reported on are those that are drawn from more than one fact.

This restriction greatly improved MYCIN's reporting in terms of its conciseness, yet it still produced valid results. When SHYSTER-MYCIN (SM-v3) restricted its reporting to conclusions made on the basis of more than one fact, it still referenced all of the provisions that the test group did.

4.2 Conciseness

The criterion that SM-v3 uses to eliminate conclusions from its report (i.e. not reporting conclusions made by applying a rule to a solitary fact) was chosen because conclusions based on a single fact are simply one-to-one mappings between one fact and another. These do not provide the user with any real information; at best, information that the user has provided is regurgitated with a slightly different wording. Conversely, a conclusion is considered interesting—worth reporting—if it is arrived at by combining several facts.

On average, SM-v3 reported only 24% of the conclusions that SM-v2 reported. It still referenced all the provisions that the test group did (as explained above), and usually only 2–3 provisions more than the test group. For each test question, SM-v2 made approximately 11 more references than did the test group.

4.3 "Correctness"

SHYSTER-MYCIN (SM-v2 and SM-v3³) agreed with the test group on each of the five questions. When opinion was divided within the test group, SHYSTER-MYCIN agreed with the majority view.⁴

5. CONCLUSIONS

SHYSTER-MYCIN successfully combines two existing expert systems. Testing of the hybrid system indicates that the approach taken to constructing the system—using rule-based reasoning to deal with statutes and case-based reasoning to deal with cases—is appropriate. A novel approach to testing allowed SHYSTER-MYCIN's reasoning, rather than its depth of knowledge, to be tested.

MYCIN was modified to operate in the legal domain, and its reporting improved by having it provide information about how it comes to its conclusions. MYCIN was further improved by having it restrict the number of conclusions that it reports on.

Linking this modified MYCIN to SHYSTER produces a legal expert system that can provide advice in an area of Australian copyright law. Because both SHYSTER and MYCIN have (separately) been shown to have general application, SHYSTER-MYCIN also has general application, and could be used in other areas of law covered by both legislation and case law.

6. REFERENCES

- [1] B. G. Buchanan and E. H. Shortliffe, editors. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley Publishing Company, 1985.
- [2] P. Norvig. *Paradigms of Artificial Intelligence Programming: Case Studies in Common Lisp*. Morgan Kaufmann, 1992.
- [3] T. A. O'Callaghan. *A hybrid legal expert system*. Honours thesis, Department of Computer Science, Faculty of Engineering and Information Technology, The Australian National University, Canberra, February 2003. <<http://cs.anu.edu.au/software/shyster/tom/thesis.pdf>>.
- [4] T. A. O'Callaghan, J. Popple, and E. McCreath. *Building and testing the SHYSTER-MYCIN hybrid legal expert system*. Technical Report TR-CS-03-01, Department of Computer Science, Faculty of Engineering and Information Technology, The Australian National University, Canberra, May 2003. <<http://cs.anu.edu.au/software/shyster/tom/tr-cs-03-01.pdf>>.
- [5] J. Popple. *A Pragmatic Legal Expert System*. Applied Legal Philosophy Series. Dartmouth, Aldershot, May 1996. ISBN 1 85521 739 2. <<http://cs.anu.edu.au/~James.Popple/publications/books/shyster.pdf>>.

³Both SM-v2 and SM-v3 operate upon the same knowledge base, so they are identical for the purposes of evaluating correctness.

⁴When the test group was divided in its opinion, it was because of differences in fact classification, rather than application of the rules. When the fact classifications made by the minority were entered into SHYSTER-MYCIN, it formed the same opinion as did the minority.