WebGrid: Knowledge Elicitation and Modeling on the World Wide Web

Shaw, M. L. G. and Gaines, B. R. (2001). WebGrid: Knowledge Elicitation and Modeling on the World Wide Web. In Industrial Knowledge Management: A Micro-level Approach, Roy, R., ed, pp. 335-348. Springer, London.

Mildred L. G. Shaw and Brian R. Gaines

Knowledge Science Institute
University of Calgary
Alberta, Canada T2N 1N4
{mildred, gaines}@cpsc.ucalgary.ca

Abstract: WebGrid-II is a knowledge acquisition and inference server on the World Wide Web that uses an extended repertory grid system for knowledge acquisition, inductive inference for knowledge modeling, and an integrated expert system shell for inference. This article describes WebGrid-II, and illustrates its use in eliciting and modeling knowledge.

1. Introduction

The development of knowledge-based systems involves knowledge acquisition from a diversity of sources often geographically distributed. The sources include books, papers, manuals, videos of expert performance, transcripts of protocols and interviews, and human and computer interaction with experts. Expert time is usually a scarce resource and experts are often only accessible at different sites, particularly in international projects. Knowledge acquisition methodologies and tools have developed to take account of these issues by using hypermedia systems to manage a large volume of heterogeneous data, interactive graphic interfaces to present knowledge models in a form understandable to experts, rapid prototyping systems to test the models in operation, and model comparison systems to draw attention to anomalous variations between experts.

However, existing knowledge acquisition tools are largely based on personal computers and graphic workstations, and their use in a distributed community involves moving software and often computers from site to site. The process of building and testing knowledge models across a distributed community would be greatly expedited if wide-area networks could be used to coordinate the activities at different sites. The initial objective of the work reported in this paper has been to use the World Wide Web to support distributed knowledge acquisition by porting existing knowledge acquisition tools to operate through the web. A further objective has been to use the web on a continuing basis to support distributed knowledge management in which the acquisition, representation and application of knowledge become an integral part of the activities of an organization.

Repertory grids based on personal construct psychology [1] have been used for knowledge acquisition since the early years of knowledge-based system development [2, 3] and have been refined over the years to support increasingly

resulting in the screen shown in Figure 2 eliciting a construct from a triad of elements.

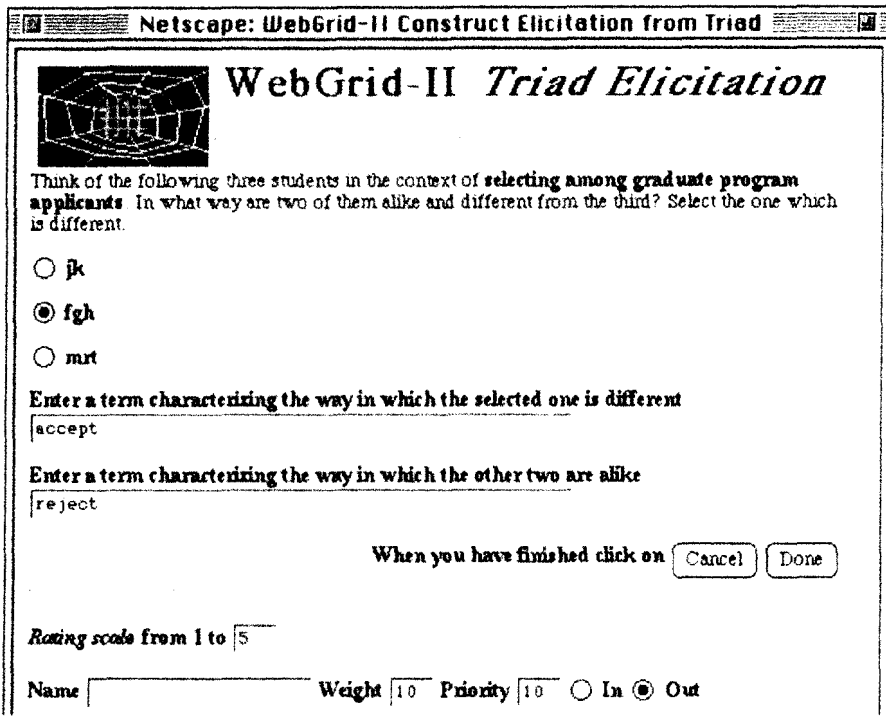

Figure 2: Construct elicitation from a triad

The user clicks on a radio button to select an element which she construes as different from the other two and enters terms characterizing the construct. The lower part of the screen is generated because advanced features were. It allows the user to vary the rating scale for an individual construct, to name the construct, to give it a weight in clustering, a priority in asking questions in an expert system application, and to specify that the construct is an input for decision making purposes or an output to be anticipated. Since the construct entered is the main decision to be made, to accept or reject an applicant, and the user wants the system to anticipate the correct decision from the way the applicant is construed she has made it an output.

When the user clicks on "Done" the server generates the screen shown on the left of Figure 3 which places a popup menu rating scale alongside each element enabling the user to rate each one along the new construct as shown on the right. She is also able to change the terms used if they seem inappropriate in the context of all the elements, and to change the ratings of the already entered elements if appropriate.

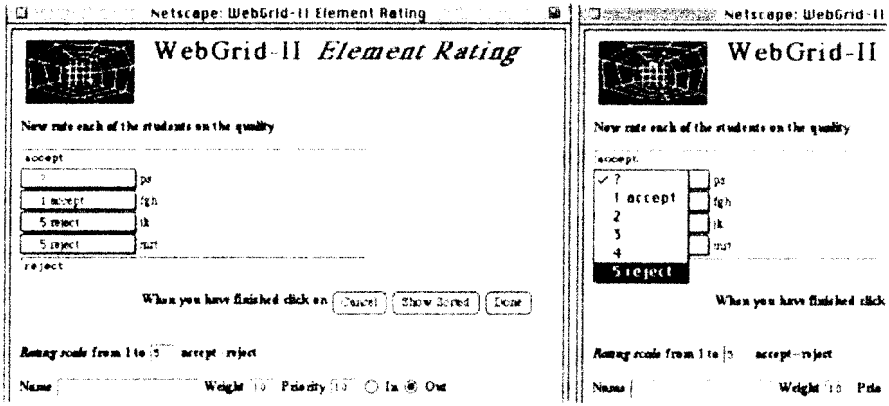

Figure 3: Rating elements on constructs

Clicking on the "Done" button in Figure 3 sends the ratings back to the server which generates the status screen shown in Figure 4.

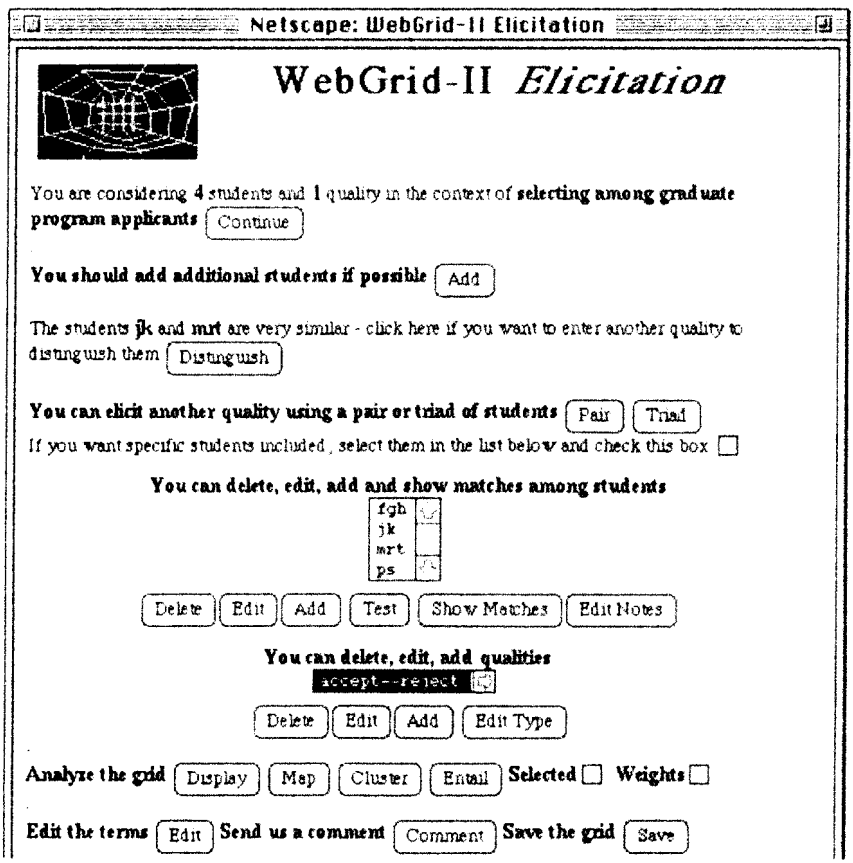

Figure 4: Status screen

This shows the elements and constructs entered, allowing them to be selected for deletion, editing and so on. It also offers various suggestions as how to continue the elicitation based on the data entered so far, facilities for analysis, saving the grid, and so on.

2.2 New data types in grids

The user decides to develop another two constructs from triads of elements and enters "high gpa—inadequate gpa" and "good references—inadequate references" since she knows that these are critical to acceptance. The applicant's gpa is a numeric rating between 0.0 and 4.0 and hence cannot be entered on a normal grid rating scale. It requires one of the new data types that we have used to extend repertory grids for knowledge acquisition.

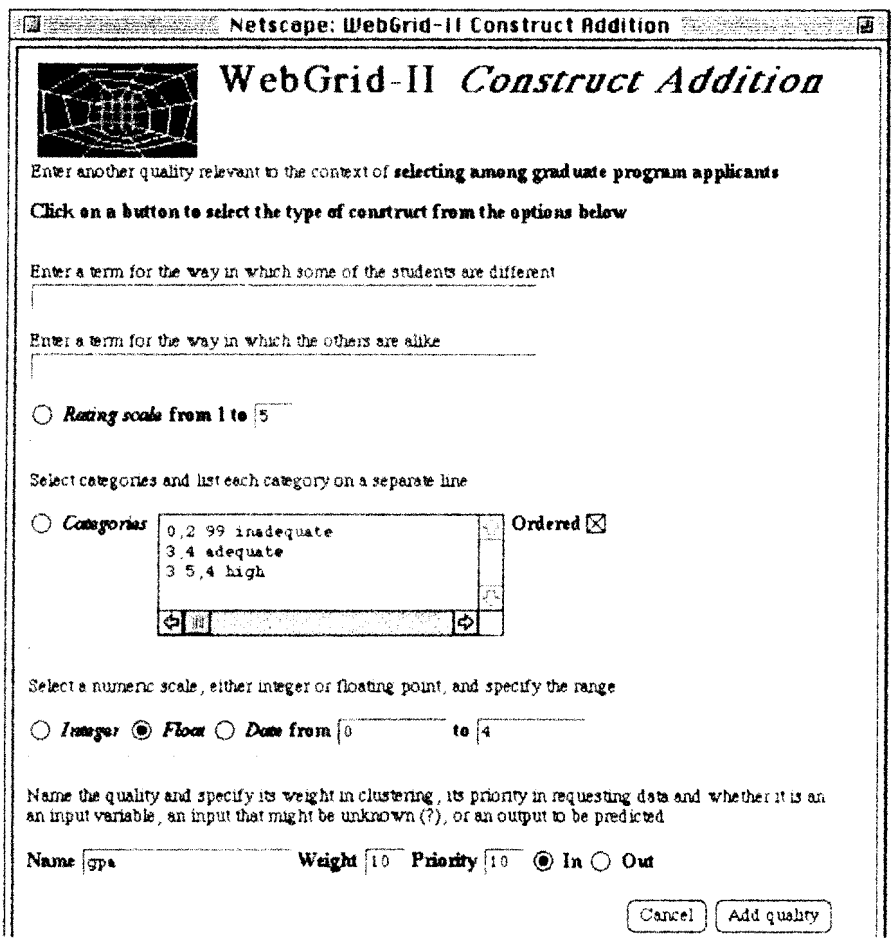

Figure 5: Adding a numeric construct with categories

Figure 5 shows the gpa categories being added as a new construct. WebGrid-II makes provision not only for rating scales but also for a set of categories, ordered or unordered, and for numerical rating scales of the type integer, float or date with prescribed ranges. The user has entered "inadequate", "adequate" and "high" as possible categories for gpa, has entered "gpa" as the name of the construct, and has noted that the value is a floating point number. The only difference the name entry makes is to shorten the process since WebGrid will generate the names "gpa = inadequate", etc, rather than have the user enter "gpa is inadequate", etc, as category names. The elements are rated on a purely categorical construct through popup menus and on a numeric construct through a numeric entry box.

2.3 Entailment analysis through rule induction

The status screen now shows three constructs and the user clicks on the "Entail" button near the bottom of Figure 4 hoping the data is sufficient to generate a model of the way in which the constructs entered entail the "accept—reject" decision. The system develops a model using the Induct algorithm [12, 13] for inductive derivation of entailments and generates the screen shown in Figure 6.

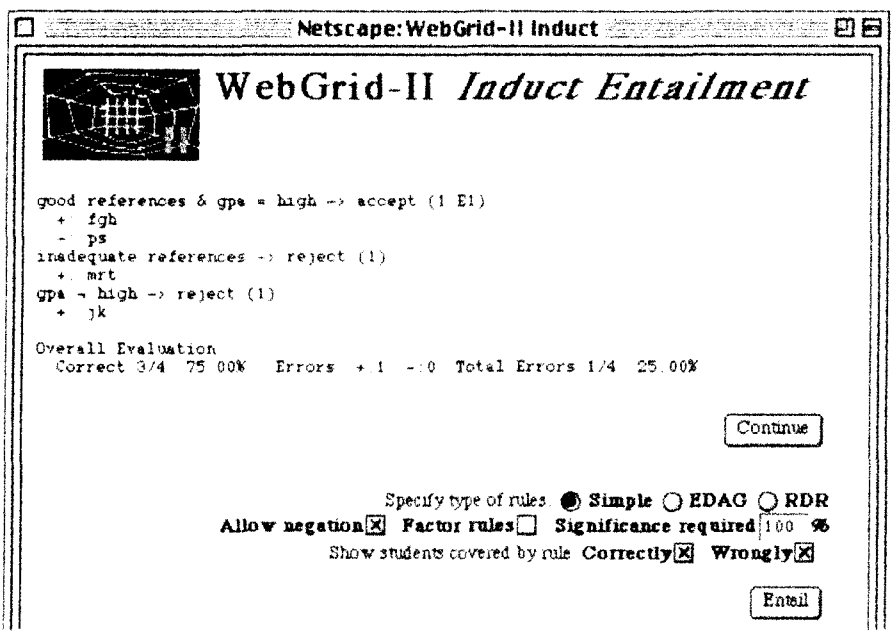

Figure 6: Rules entailed by the grid after 4 elements and 3 constructs

She is disappointed to receive the entailments shown which look sensible but work for only 3 cases in 4. The output consists of a set of rules defining the anticipatory entailments and each rule is followed by the positive elements where the correct decision is anticipated and the negative ones where that anticipated is incorrect. On thinking about it she realizes that the applicant "ps" has a high gpa and good references but was rejected because there was no-one willing to

supervise him. She returns to the status screen, adds a construct "willing supervisor—no supervisor" and clicks on "Entail" again. This time the rules entailed by the grid returned as shown in Figure 7 correctly anticipate the decisions made for all 4 cases.

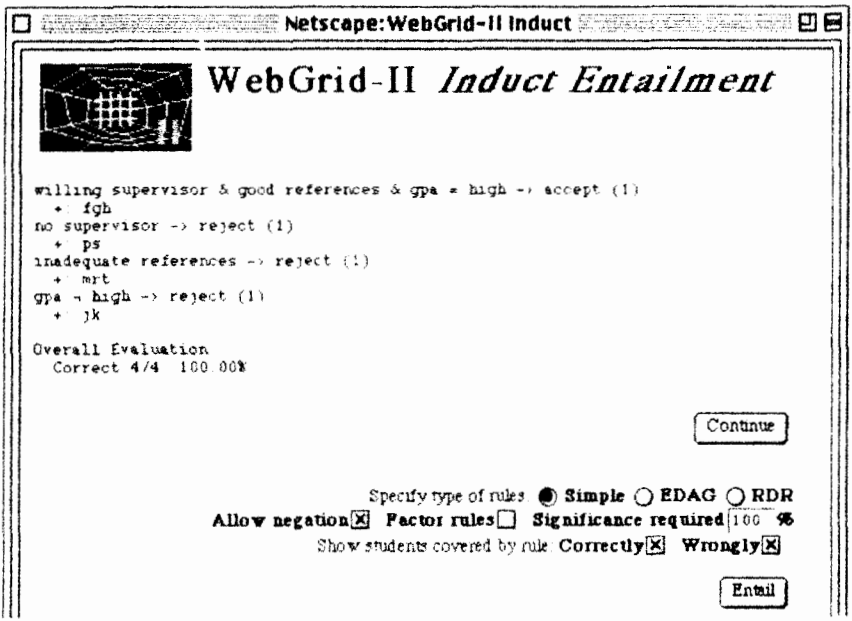

Figure 7: Rules entailed by the grid after 4 elements and 4 constructs

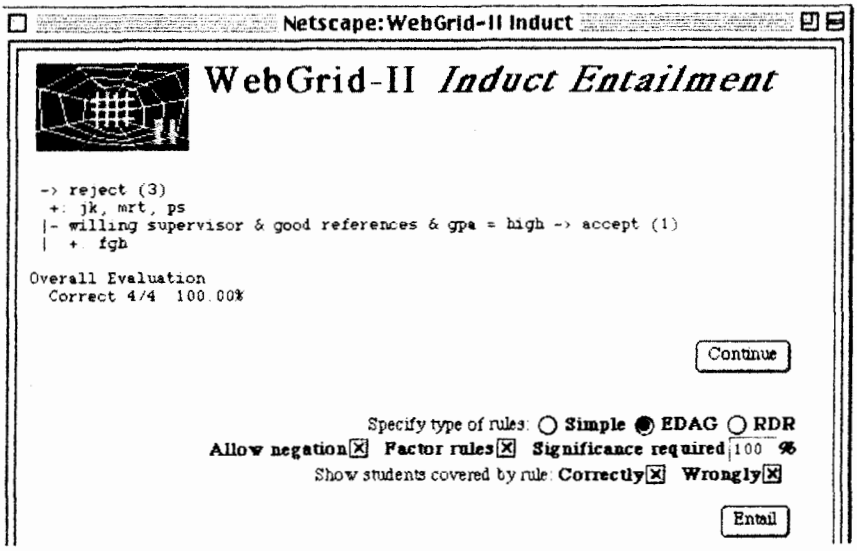

Figure 8: Rules with exceptions entailed by the grid after 4 elements and 4 constructs

The user knows that Induct is also capable of producing default rules with exceptions [14, 15] that are often more natural in their correspondence to psychological models, and runs Entail again with "EDAG" selected in Figure 7. The system generates the screen of Figure 8 which has the default rule to reject all applicants unless they have a high gpa, good references and a willing supervisor—a good initial characterization of the selection process.

2.4 Inference through an embedded expert system shell

The user decides to test the entailed rules on further cases and returns to the status screen of Figure 4 which now shows four constructs. She clicks on the "Test" button under the list of elements and the system returns the screen shown in Figure 9 which allows her to specify a test case and click on "Infer" to see what decision is anticipated.

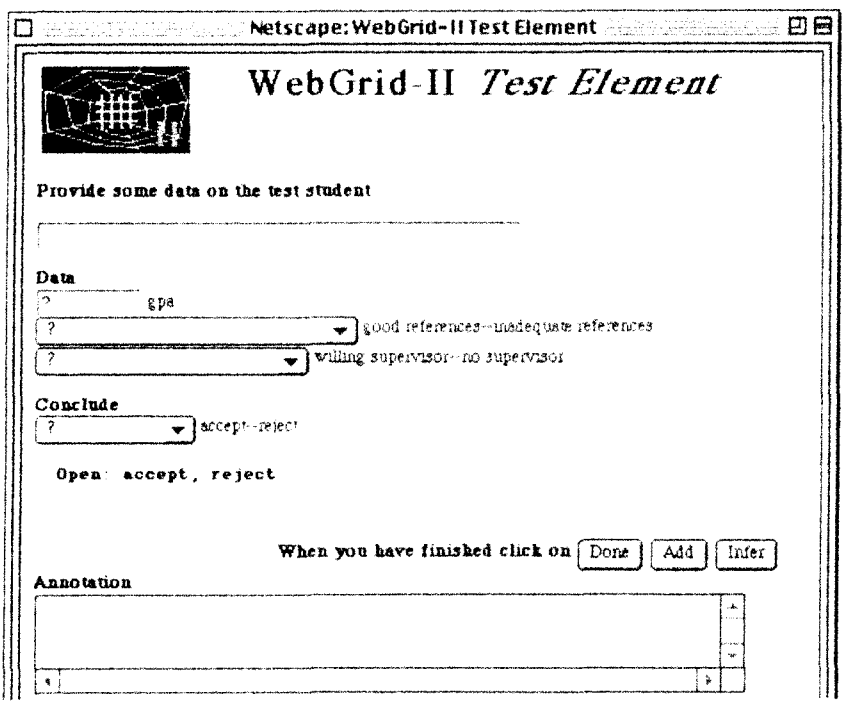

Figure 9: Entering a test case into the entailed expert system

The user does this for several similar cases and the conclusion returned is correct but when she does so for the applicant "lemp" as shown in Figure 10 the conclusion is "reject" when she knows the applicant has been accepted.

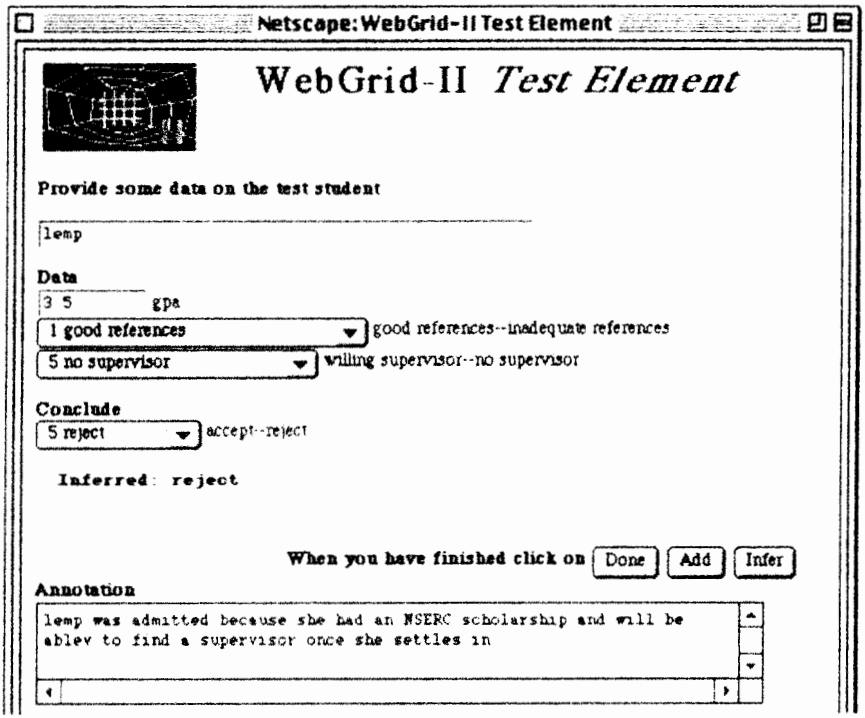

Figure 10: Conclusion anticipating the decision for a test case

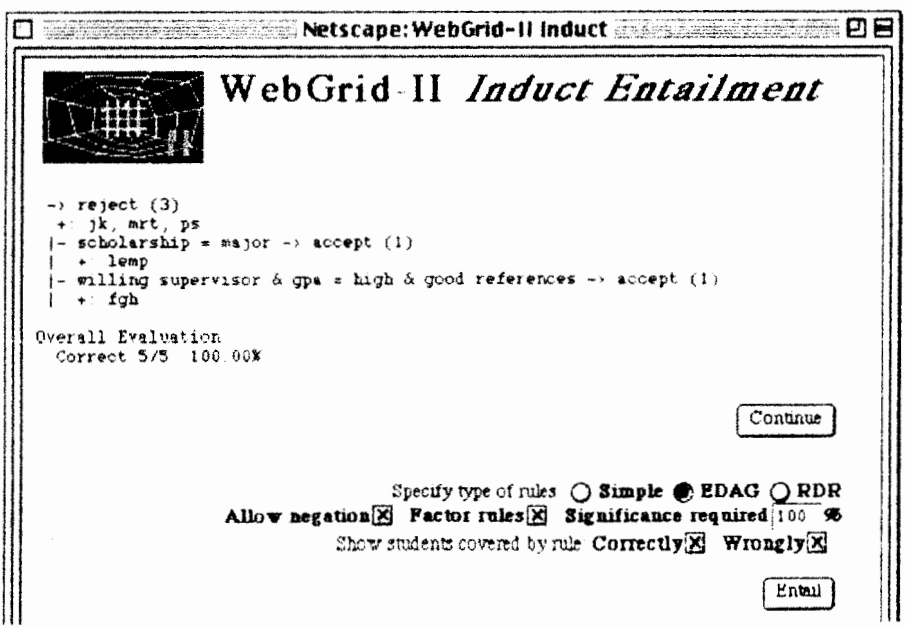

Figure 11: Rules with exceptions entailed by the grid after 5 elements and 5 constructs

The user realizes that applicant "lemp" was admitted because she had a major scholarship, corrects the conclusion to "admit", and clicks on "Add" to add this applicant as a new element in the grid. The system returns to the status screen of Figure 4 with "lemp" added and the user clicks on "Add" under the construct list to enter the categorical "scholarship" is "major" or "none" She then clicks on "Entail" and the system generates the entailed rules with exceptions as shown in Figure 11 which now account for "lemp" also. The default rule is to reject all applicants unless they have a major scholarship or a willing supervisor, high gpa and good references when they should be accepted unless they have no supervisor and no scholarship—a more refined characterization of the selection process. However, she also realizes there are applicants with major scholarships who have not been admitted because they have a poor gpa or inadequate references.

The user enters some relevant examples, returns to test case entry again and this time come across a applicant who was admitted to a new program, an industry-based software engineering degree with the requirement that the applicant have substantial industry background. She enters some relevant cases and constructs leading to the entailed decision rule shown in Figure 12, that applicants should be rejected by default unless they have good references and in addition either have a high gpa but not no scholarship and no supervisor, or are a software engineering application with industry experience and an adequate gpa.

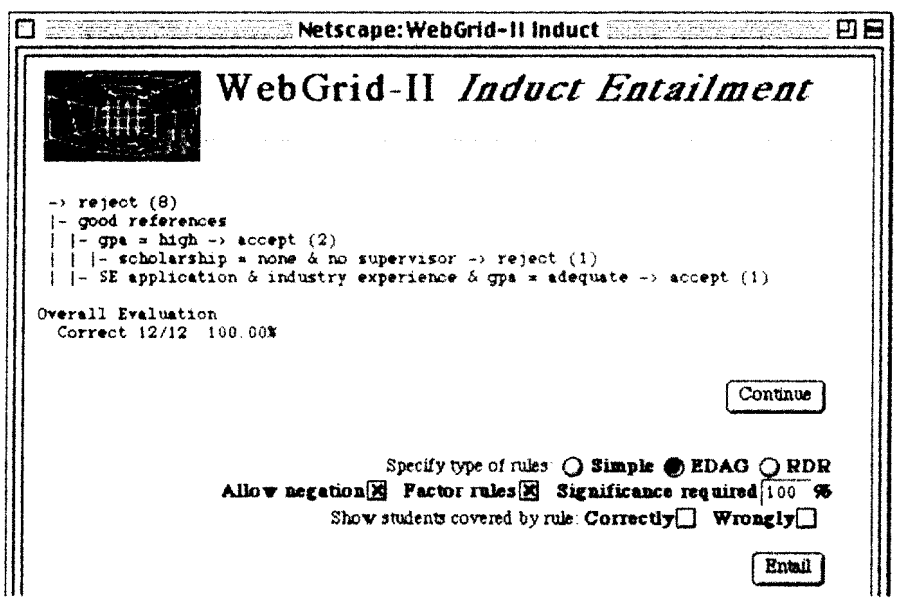

Figure 12: Rules with exceptions entailed by the grid after 12 elements and 7 constructs

Note how the decision rule entailed by the grid is becoming more refined as it encompasses more cases and also how the Induct algorithm generating it changes its structure to be as comprehensible as possible as it grows in complexity. This attempt to generate an anticipatory structure which is natural for people is an important feature of EDAG generation [15]. It is also noteworthy that a fairly

small grid of only 12 elements and 7 constructs can be used to generate such a complex decision rule. This is because the repertory grid elicitation technique with PEGASUS-style [16] feedback leads to the elicitation of typical elements and relevant constructs, allowing knowledge to be elicited highly efficiently compared with rule induction from a database where thousands of examples might be required to derive the same knowledge [12].

3. Conclusions

The technologies, tools and protocols developed for the World Wide Web make it possible to provide knowledge acquisition, representation and inference systems internationally to anyone with access to the Internet. Such knowledge-based technologies integrate naturally with other applications and hypermedia systems to provide a distributed knowledge medium capable of supporting the knowledge processes of professional communities.

This paper has illustrated what is possible through examples of the development of knowledge modeling tools operating through the web to provide knowledge acquisition, representation and inference through repertory grids. These systems were developed primarily for locally managed projects, and their availability has only been advertised through specialist list servers for the knowledge acquisition and personal construct psychology communities. However, in the period June 1995 to June 1997 WebGrid was accessed from over 5,000 different sites in some 60 countries, delivering over 200Mbytes of grid data in some 70,000 transactions. The web is essentially an anonymous medium and we do not track the activities of outside users. However, occasionally users contact us to discuss their applications and, if the local network goes down, we receive mail from serious users reporting the problem and requesting notification of when the knowledge acquisition programs will be available again. An interesting example of serendipitous use was a masters student in Holland studying operators' models of nuclear reactors who came across the grid elicitation system on the web in September 1996 used it with his experts to elicit their conceptual models, and included these in an additional chapter in his thesis for examination in October 1996—the web certainly accelerates research processes.

In conclusion, we note that the results presented in this paper demonstrate that the World Wide Web client server model operating through the Internet can effectively support the porting of existing knowledge acquisition, modeling and representation tools to operate in a distributed, open architecture environment. This opens up new possibilities for collaborative research in which the methodologies and tools developed by one group are made available on the network for use by other research groups world-wide. It enables past studies of the integration of hypermedia, knowledge acquisition and inference tools within a local environment [17, 18] to be generalized to integration on a distributed basis across wide area networks. A major advantage of this approach is that tools can be operated at the site of the originators responsible for them, and their maintenance, upgrading and monitoring can be done very much more effectively than when the software is distributed to other sites.

We envision a collaborative research program internationally in which the first objective is to make tools available across the Internet in interoperable forms, and

the second objective is to develop some integrated knowledge management systems that use these tools in a variety of configurations to provide a new generation of knowledge-based systems

Acknowledgments

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada.

URLs

WebGrid can be accessed at http://tiger.cpsc.ucalgary.ca/WebGrid/
Related papers on WebGrid and World Wide Web can be accessed through http://www.RepGrid.com/articles

References

1. Kelly, G.A., The Psychology of Personal Constructs. 1955, New York: Norton.

2. Shaw, M.L.G. and B.R. Gaines, A computer aid to knowledge engineering, in Proceedings of British Computer Society Conference on Expert Systems. 1983, British Computer Society: Cambridge. p. 263-271.

3. Boose, J.H., Personal construct theory and the transfer of human expertise, in Proceedings AAAI-84. 1984, American Association for Artificial Intelligence: California. p. 27-33.

4. Boose, J.H., et al., Knowledge acquisition techniques for• group decisions support. Knowledge Acquisition, 1993. 5(4): p. 405-447.

5. Gaines, B.R. and M.L.G. Shaw, Basing knowledge acquisition tools in personal construct psychology. Knowledge Engineering Review, 1993. 8(1): p. 49-85.

6. Ford, K.M., et al., Knowledge acquisition as a constructive modeling activity. International Journal of Intelligent Systems, 1993. 8(1): p. 9-32.

7. Shaw, M.L.G. and B.R. Gaines, Comparing constructions through the web, in Proceedings of CSCL95: Computer Support for Collaborative Learning, J.L. Schnase and E.L. Cunnius, Editors. 1995, Lawrence Erlbaum: Mahwah, New Jersey. p. 300-307.

8. CPCS, RepGrid 2 Manual. 1993: Centre for Person Computer Studies, PO Box 280, Cobble Hill, BC, Canada V0R 1L0.

9. Shaw, M.L.G., B.R. Gaines, and M. Linster, Supporting the knowledge engineering life cycle, in Research and Development in Expert Systems XI, M.A. Bramer and A.L. Macintosh, Editors. 1994, SGES Publications: Oxford. p. 73-86.

10. Shaw, M.L.G. and B.R. Gaines, Group knowledge elicitation over networks, in Research and Development in Expert Systems X. Proceedings of British Computer Society Expert Systems Conference, M.A. Bramer and A.L. Macintosh, Editors. 1993, Cambridge University Press: Cambridge, UK. p. 43-62.

11. Gaines, B.R. and M.L.G. Shaw, Eliciting knowledge and transferring it effectively to a knowledge-based systems. IEEE Transactions on Knowledge and Data Engineering, 1993. 5(1): p. 4-14.

12. Gaines, B.R., An ounce of knowledge is worth a ton of data: quantitative studies of the trade-off between expertise and data based on statistically well-founded empirical induction, in Proceedings of the Sixth International Workshop on Machine Learning. 1989, Morgan Kaufmann: San Mateo, California. p. 156-159.

13. Gaines, B.R., Refining induction into knowledge, in Proceedings of the AAAI Workshop on Knowledge Discovery in Databases. 1991, AAAI: Menlo Park, California. p. 1-10.

14. Gaines, B.R., Exception DAGs as knowledge structures, in AAAI-94 Workshop: KDD'94 Knowledge Discovery in Databases. 1994, AAAI: Menlo Park, California.

15. Gaines, B.R., Transforming rules and trees into comprehensible knowledge structures, in Knowledge Discovery in Databases II, U.M. Fayyad, et al., Editors. 1995, AAAI/MIT Press: Cambridge, Massachusetts. p. 205-226.

16. Shaw, M.L.G., On Becoming A Personal Scientist: Interactive Computer Elicitation of Personal Models Of The World. 1980, London: Academic Press.

17. Gaines, B.R. and M. Linster, Integrating a knowledge acquisition tool, an expert system shell and a hypermedia system. International Journal of Expert Systems Research and Applications, 1990. 3(2): p. 105-129.

18. Gaines, B.R., A. Rappaport, and M.L.G. Shaw, Combining paradigms in knowledge engineering. Data and Knowledge Engineering, 1992. 9: p. 1-18.