Artificial Intelligence and Law **3:** 217–220, 1995.

Book Review

Robert F. Walker. *An Expert System Architecture for Heterogeneous Domains*. Ph.D. thesis, Vrije University, Amsterdam. Zeist: A-D Druk, 1992. 385 pp.

The PROLEXS project has been a significant enterprise in the field of AI and law. It began in 1986 and has involved a team of people at the Vrije University in Amsterdam. It is perhaps best known through papers in the various international conferences on AI and law (Oskamp *et al.* 1989, Opdorp *et al.* 1991) and through an article in the *International Journal of Man-Machine Studies* (Walker *et al.* 1991).

The target system was one which would give advice to tenants on landlord and tenant law in the Netherlands. It can thus been seen as at once pragmatic, in that there was a clear focus of domain and intended user, and as theoretical, in that it was undertaken by a team of academics with some rather clear ideas as to how legal knowledge-based systems should be built. The key idea was that since the relevant legal knowledge needed to be drawn from a range of heterogeneous sources, such as legislation, case law, experience and the like, the representation of this heterogeneous material should use a variety of knowledge representation formalisms, so that the formalism most appropriate to the source could be used. Whilst this idea is not unique, since many people have argued for and against the suitability of various formalisms with respect to various types of source material, PROLEXS is unusual in the extent to which the idea was carried out in the context of a practically focused project. To build a coherent system it was necessary to integrate the different representations. Rob Walker was a key computer science-oriented member of the team and was largely responsible for the design of the underlying blackboard architecture which was used to effect the integration. This thesis describes that aspect of the work, in the form of a general system called EXPANDER, but throughout the real focus is on the system which underlay the PROLEXS project.

Part I offers a short introduction and Part II some general analysis of the problem domain and the various representation and reasoning paradigms that motivate the desire for a heterogeneous system. The bulk of the book is in Part III, which provides an admirably full description of the architecture to realise such a system. Part IV describes the application of this architecture to the legal domain, and includes some interesting discussion of legal knowledge representation issues. Particularly useful is the discussion of neural networks and their potential for application in law: neural networks became popular during the course of the project, and it is an illustration of the flexibility of the architecture that they could be incorporated within the design. Part V provides some additional implementation details and Part VI some – perhaps too brief – evaluation and conclusions.

Anyone who has seen the papers on the PROLEXS system and wants to find out more about the project will want to read this book, which supplies all the details that are inevitably left out of conference and journal papers. The technically minded will find all the design and implementation detail they could hope for, and it is interesting both as an account of the problems encountered and of the particular solutions found in the course of the project. Those of a less technical bent will still find Parts II and IV of interest, as they provide a good contribution to the debate as to whether heterogeneous sources require heterogeneous representations, and the questions that must be asked if heterogeneous representations are to be combined in a system.

Walker identifies three problems that need to be addressed: how to coordinate the heterogeneous knowledge sources; how to resolve conflicts between them; and how to integrate their contributions into an intelligible explanation (p. 65). Coordination is effected by EXPANDER broadcasting messages to and receiving messages from the various knowledge sources. These messages will be expressed in EXPANDER's language (L), and it may be necessary for them to be translated into a form understandable internally by the knowledge sources. The translation has to be the responsibility of the knowledge sources, since all their internal details are hidden from EXPANDER. This requirement does, however, mean that considerable coordination between the knowledge groups and L, the language used by EXPANDER, is necessary. Since knowledge groups can only get as input and offer as output messages in L, L must be rich enough to express all the input required by the knowledge groups, and the statements derived by the knowledge groups must be expressible in L if they are to contribute to the overall problem-solving process. Such commonality is unlikely to happen by chance: clearly the authors of L must be aware of the input terms required by the knowledge groups and the authors of the knowledge groups must be aware of the predicates required by other knowledge groups so that their conclusions can be used. L is therefore by no means as independent of the knowledge groups as one might have hoped. Walker indeed addresses this problem, but seems to me rather optimistic in his contention that "as long as a team [constructing a knowledge group] broadcasts any symbol they create and associate a well-defined meaning to the symbol their communication with other teams can be fairly limited" (p. 274). I would expect a great deal of communication to be needed, and that the needs of other knowledge groups would have a significant impact on the design of any given knowledge group.

Another key problem of coordination is the use of a certainty factor to be associated with the various statements posted to the blackboard. It may be that different knowledge groups will treat uncertainty differently: one might need probabilities, another subjective certainty factors à la MYCIN, a third fuzzy logic, whilst the fourth could be a neural net. All will associate a number between 0 and 1 with a proposition, but the interpretations to be placed on these numbers cannot be reconciled. Again Walker acknowledges this difficulty, and says that "it is up to the knowledge group developers to agree on a method of certainty measure assignment that can be meaningfully compared across knowledge groups" (p. 274). Certainly this proposal would solve the difficulties, but it is an important concession to homogeneity, made in an area where the advantages of heterogeneity

are perhaps clearest. The treatment of uncertainty was not a problem in the PROLEXS system, which did not use quantified uncertainty, but could be critical in some applications.

These and like considerations suggest that the various knowledge groups which will form the application need to be designed together, for the specific application. Adding a new knowledge group will require consideration of whether changes are required in the existing knowledge groups, and could lead to the need for substantial alteration. Importing a knowledge group designed by a different team is likely to require extensive changes, both to the imported and existing knowledge groups. This is a little disappointing since one of the hopes raised by the architecture was that an application might be constructed by putting together some preexisting knowledge groups, perhaps with some specially written ones, without the need to modify them for the new context. This straightforward reuse is an important goal in software engineering, but it requires looser module coupling and more data hiding than appears supported in PROLEXS. Walker himself sees knowledge reuse as only a limited possibility, and sees instead the major reuse opportunities as applying to the knowledge representation formalisms and inference engines of the knowledge groups.

Turning next to the resolution of conflicts, further problems arise. If two knowledge sources post contradictory statements on the blackboard, EXPANDER must resolve the conflict. Walker gives a mechanism for conflict resolution, but it is not, in my view, entirely satisfactory. The problem is that the precise reasoning underlying a statement is not available to EXPANDER, since the implementation details of the knowledge groups are hidden, and thus it can only go on very general information such as the knowledge group which supplied the statement or the order in which the statements were placed on the blackboard. Whilst it might be possible to rank knowledge groups (e.g. perhaps case law defeats legislation), it is much more likely that we would want to consider the arguments as well, so that, for example, a good case law match may defeat legislation, but legislation is to be preferred over a weak match. Resolving conflicts is difficult at the best of times, but when the conflicting statements are derived from different representations, using different reasoning methods, and where the precise line of reasoning is unavailable, then a reliable general solution seems to be impossible. Of course, EXPANDER does provide a mechanism which will resolve conflicts, but one suspects that some fairly good fine tuning of the knowledge groups would be needed to make it perform in an acceptable fashion.

Taking the difficulties with coordination and conflict together, it seems to me that the problems of heterogeneity which are avoided by providing the possibility of different styles of representation resurface both in coordinating the design of the different knowledge groups and of getting the appropriate behaviour when they interact. The heterogeneous nature of sources needed by legal knowledge-based systems does present a set of challenges, but allowing the use of different knowledge representations may only defer them, not remove them. Despite the attractions of heterogeneous representation, well argued for in this book, I remain unconvinced that the difficulties are not best solved during the process of representing the disparate knowledge sources in a single formalism.

The third problem raised by Walker relates to combining the various contributions into an intelligible explanation. This issue is important, and the book contains a good discussion of how it was done in PROLEXS, but I shall not discuss it further in this review.

None of what I have said is intended to detract from the PROLEXS project or from Walker's book. It is only through such a thorough investigation of the approach that a context in which these issues can be sensibly debated is created. Both the project and the book make a significant contribution to the debate of how best to represent legal knowledge.

T. BENCH-CAPON

University of Liverpool

Acknowledgement

I would like to thank Anne Gardner for her helpful comments on an earlier draft of this review.

References

Opdorp, G. J. van, Walker, R. F., Schrickx, J. A., Groendijk, C., & Berg, P. H. van den. 1991. Networks at Work: A Connectionist Approach to Non-Deductive Legal Reasoning. In Proceedings of *The Third International Conference on Artificial Intelligence and Law*, Oxford, 278–287. New York: Association for Computing Machinery.

Oskamp, A., Walker, R. F., Schrickx, J. A., & Berg, P. H. van den. 1989. PROLEXS, Divide and Rule: A Legal Application. In Proceedings of *The Second International Conference on Artificial Intelligence and Law*, Vancouver, 54–62. New York: Association for Computing Machinery.

Walker, R. F., Oskamp, A., Schrickx, J. A., Opdorp, G. J. van, & Berg, P. H. van den. 1991. PROLEXS: Creating Law and Order in an Heterogeneous Domain. *International Journal of Man-Machine Studies* 35: 35–67.