"The Röle of AI in Requirements Engineering"

Position Paper for the
Dagstuhl
Seminar on
"System Requirements: Analysis, Management and Exploitation"
October 1994

Kevin Ryan
University of Limerick
IRELAND
ryank@ul.ie

1. Summary

There are two orthogonal aspects to the interaction between AI and RE. The more obvious is that we may look to AI to help solve many of the most difficult and crucial problems facing RE particularly those concerning knowledge storage and user acceptance. After a brief summary of the relationship between the fields of AI and RE some typical such problems are identified and the potential contribution of AI is estimated. The second aspect is that the achievements, limitations and methodology of AI have lessons for the RE community. AI's limited success in solving many of its problems should form a cautionary reminder to RE researchers. We must be aware that realism and practicality are as important to effective deployment as are grand plans or novel concepts. The distinguishing characteristics of successful and accepted AI-based systems are used to identify some possibilities and pitfalls in the future.

2: AI & RE

No doubt many of the participants at this workshop have extensive knowledge of some or all of the AI field. For others, it may be worthwhile to state that it is difficult to characterise the field of Artificial Intelligence both because of its breadth and its dynamic nature. AI has grown from being an esoteric and often incredible niche to embrace many subdisciplines including, "but not limited to" Planning, Game Theory, Natural Language Processing, Machine Learning, Neural Networks, Knowledge Representation and Expert Systems. At this stage many of these areas have their own distinct research communities supported by specialist journals and conferences. There is wide variation however in the extent to which the results of these field have been accepted in to everyday use.

The much more recent field of RE includes (approximately) all that is concerned with the elicitation, specification and analysis of requirements for software intensive systems. Most of us would see RE as a subdiscipline of software engineering - witness the ICSE award to Anthony Finklestein and Steve Fickas - but others would argue that RE is essentially interdisciplinary or that it must be subsumed (along with software engineering?) into the (re-) emerging discipline of Systems Engineerng.

3. Possible AI applications in RE

There are many appealing possibilities for the application of AI approaches within the field of RE. Many of these are being investigated at present and will be presented at this workshop. Some of the more obvious examples are :

- Natural Language Processing technology could be used in the capture, paraphrasing or comparison of requirements.
- Game Theory might be useful in for evaluating alternative requirements, especially in representing the trade-offs between stakeholders' priorities.
- Plan Generation, most often used for synthesis problems in AI, could be applied to the systematic refinement or revision of requirements.
- Machine Learning, with its rapidly growing subfield of neural networks, is a form of heuristic model building that might be acceptable in some domains.
- Expert Systems, where the reasoning strategy is explicitly represented, would be more suitable for other domains and might also be used to advise on modelling choices. • The subfield of AI known as Knowledge Elicitation clearly has much in common with the area of requirements capture. A recent paper by Shaw and Gaines [RESS94] illustrates this overlap very convincingly and argues strongly for more sharing of technologies and paradigms between the two research groups.
- Last, but by no means least, it is evident that Knowledge Representation issues, such as tractability, expressive power, interchange formats.and efficient retrieval, are fundamental to most if not all areas of RE.

No doubt there are many more possibilities but even from this limited list it should be clear that these possibilities vary both in their practicality and in their potential payback. It might be useful to ask why AI seems so widely applicable before going on to

4. Dealing with Uncertainty

Shaw and Gaines state that Knowledge Based systems are best applied to uncertain domains. There is no doubt that RE involves considerable uncertainty both in what is being dealt with ("product uncertainty") and in how to proceed ("process uncertainty"). In considering the "informal-formal" transition that is central to the RE process,

Reubenstein [ ] identifies six causes of uncertainty. These are abbreviation, ambiguity, poor ordering of information, contradiction, incompleteness and inaccuracy. We can consider how AI might help us deal with each of these. Knowledge rich domain models (NATURE, ARIES) can help us address both abbreviation ("The client assumes we know or understand") and - to a limited extent - incompleteness (" the client forgets to tell us"). The critiquing approach, used with some success in medical expert systems, could deal with some forms of ambiguity ("the client assumes we understand correctly") and inaccuracy ("the client makes a mistake"). Contradiction ("client rules have exceptions") is handled in knowledge-based systems by truth mainenance systems. The difficulties posed by poor ordering of information ("the client is not a teacher") are a major part of the "information bottleneck" and may be tackled using all of the many techniques of knowledge elicitation. At this point RE overlaps into cognitive and social science and must be congniscent of the strengths and limitations of these "softer" disciplines. However all of the approaches suggested here can help increase the clients' trust in the eventual system by reducing the level of uncertainty in its requirements definition.

5. Lessons from AI

But some would argue that bright, young, innnocent RE should be wary of involvement with tired, old, discredited AI. After all, the field of AI has not had a great press with the general public or with the computing community in general. This may be attributed to a number of causes including, as is often pointed out, the tendency to say that once a problem has been solved it is no longer part of AI. But there is no doubt that at least part of the cause has been the underestimation of the difficulties involved. Allied to occasional fits of hubris, this has the unfortunate effect that many of the signiificant achievements of AI are ignored or underestimated outside that field. Achievements such as the pervasive use of expert systems, heuristic search algorithms, (restricted) NLP tools and adaptive systems based on neural networks. Not all areas have been equally successful however - we still await the first industrial strength voice recognition system - so RE might benefit from a quick look at what characterises the more successful ones.

Four characteristics can be identified. Firstly, successful AI has generally been within narrow niches or sub-domains. Generalised intelligence or common sense has proven much more elusive than specialised, tightly constrained "knowledge". No doubt this reflects our continuing struggle to understand our own higher, and hence more abstract, functions. Secondly, we can see that it has proven easier, particularly in expert systems, to provide passive rather than active assistance. The critiquing approach mentioned earlier is an example of this but similar conclusions can be

drawn from experience with knowledge based CASE tools. Users are much happier to dialog with a "helpful advisor" than to put their trust in any mysterious oracle. The third characteristic is related to the second. AI in general, and knowledge-based systems in particular, work best when suboptimal solutions are acceptable, what Simon termed "satisficing". Furthermore, since "knowledge" according to Shaw and Gaines "is more than opinion but less than fact", we must be willing to accept occasional wrong answers from our knowledge-based systems. This is a crucial issue for RE researchers, if we are to use KB technology successfully. Finally, KB systems are only acceptable to users when they can provide, to the level of detail needed, explanations of what they do and why they do it. While some systems, such as adaptive modems, need provide no explanation, most must re-assure the responsible human user that the decisions being made are consistent with the facts and, in general, correct. The level of explanation required in RE is likely to be very high. Only full scale trials, on real applications, can help determine whether AI-based RE systems can win the trust of the users currently responsible for requirements definition. Experience with KB-systems is that such trials are expensive, difficult and very rewarding.

Conclusions

It seems clear that, despite its shortcomings, AI has much to contribute to the RE field. It is particularly useful as an example of a technology that has won hard-earned user acceptance, albeit in limited nices to date. I think it safe to say however that it will be some time before we have industrial strength AI systems for the RE process.