

Integrating Research into the Core Curriculum - A Case Study

Erika Rogers

This presentation describes an approach to providing undergraduates with a "real-world" research experience within the context of a core computer science course. As a required course, Introduction to Artificial Intelligence (AI) is typically taught in the junior/senior year, and normally covers fundamental topics in AI as well as Lisp programming. While this material is in and of itself quite challenging, I wanted to give the students an experience of what is involved in a real AI research project, such as the Visual Interaction Assistant for Radiology [1], which has been under development over a number of years. This intelligent assistant is based on extensive cognitive data collection from radiologists reading chest x-rays, and there are several open analysis problems which still need to be resolved. It seemed that the "think-aloud" reports of the radiologists could be used as a vehicle to help the students learn how to extract computationally meaningful information from a domain expert's thoughts and words. Many issues arising in this exercise are applicable to the development of expert systems in other domains as well, for example: handling specialized vocabulary, developing quantitative measures and qualitative concept labeling procedures, applying subjective ratings, aggregating data, and interpreting results of the analyses both individually and in the context of the whole experiment.

The required tasks included a time analysis using the original video data, a concept analysis where the words in the protocols were categorized and labeled, and a performance analysis where the subjects' diagnoses were compared with those of the project expert. The students were also required to prepare a final project report on their results, with the help of a detailed outline, which contained suggested section headings and questions to be addressed in each section. This was in keeping with the current trend of "writing across the curriculum", since typically, our computer science students are not taught how to prepare research-style papers. This guidance helped them develop a methodology for pulling together their ideas

and presenting them in a coherent manner.

There were several positive outcomes of this exercise: a large number of the project writeups were extremely well-done – in the students' grasp of what the VIARAD project was all about, what their own results meant when aggregated across subjects and across cases, and the general quality of writing (which was much higher than the typical standard for other computer science students I've seen). Some of their analyses were insightful, and their summaries of what they learned from doing this project were surprisingly consistent - the majority of students claimed that they went from total confusion in the beginning to "seeing the light" by the concept analysis stage. Most felt that this experience greatly enriched their understanding of the scope of fundamental AI principles, much more so than the toy problems which are typically used for illustrative purposes. A concrete outcome was that many of the students expressed an interest in further pursuing this type of research, and 8/23 (35% of the class) completed a Research Interest Application form asking to work on VIARAD or some related project in the future. Thus the experiment was far more successful than had been anticipated, and many of the benefits are still to be reaped. The results which the students produced are actually new to the project, and once the figures have been verified, can be incorporated into a publication. This experience will be reviewed from the point of view of "lessons learned", and refinements on the initial design will be discussed.

References

- [1] Rogers, E., 1995. "VIA-RAD: A blackboard-based system for diagnostic radiology". *Artificial Intelligence in Medicine* 7, pp. 343-360.