Teaching Students with Practical Concerns

Jonathan Hodgson
Department of Mathemaitcs and Computer Science
Saint Joseph's University
Philadelphia. PA 19131
jhodgson@sju.edu

I view my primary goal as a teacher of AI to be to try and give a coherent view of what is by its very nature an incoherent field. I think that for many students the first encounter with AI is unsatisfactory because they are unable to build their own model of what goes on in AI. They are exposed to different techniques for search, for knowledge representation and inference which they do not get a chance to integrate.

One of the difficulties inherent in the presentation of AI is that the discipline is rife with internal divisions; neat versus scruffy AI, artificial versus intelligence, symbolic representations versus neural networks. These differences are critical because they can determine both the form that an AI application takes and the approach to the subject as a whole. Many of these differences are philosophical in nature in that they concern what one thinks AI should be about. Students can easily be confused when both sides of a schism are presented, with the result that they abandon the field. The problem is to find an honest way to own up to the divisions without fragmenting the course.

I take the position that the primary task of AI is to build systems that are capable of performing tasks that are not well posed. This does not mean that systems should be justified solely by what they do. It is important to understand how they work. This can best be done for systems that are neatly put together. The systems will be large and complex, but they will contain components that are smaller, simpler and in some cases well defined.

Most of the students that I teach are impatient with theory, they grasp ideas best when they can see them in practice. If they see a program that implements a simple theorem prover that helps them with the idea of a theorem prover in general. For this kind of student I would like to have a largish system that is well structured and accessible so that it can be modified. I want the system to be written in Prolog (here is another "religious" dichotomy that I feel no need to apologize for.) With such a system one could discuss the options that are available for each part of the system, perhaps even assign students to implement one of the alternative options. If the system has enough parts the instructor would be able to choose those parts that corresponded to his/her preferences in AI.

It is in the nature of AI that there is no one system that will be an acceptable model for the community at large. There are just too many applications out there, each with its own demands. It is unrealistic to expect the system to do all the things we would like it to do, ML, vision, NLP, motion, planning etc.. On the other hand it is not unreasonable to ask for a system that does some of these things and could be enhanced to do others, so that one could discuss the options with students.

Since writing the original version of this position paper I have given some more thought to the question of motivating practically oriented students. The ones that I see have one eye firmly on the job market. They read the want-ads. These ask for people with database or networking experience. Good examples of AI applications in these fields will help motivate these students. Thus I would extend my desiderata to require the system to have an obvious connection to either networking or database.

With this in mind I have been working on an example prompted by an application that out career services people have requested. They would like to have an online database of job files. Students looking for a job ask to see job opportunities that fit a relatively simple set of criteria. With the files on line there is an opportunity for an AI application that tags the files so that they can be retrieved according to these criteria. I hope to have enough of this portion of the project written by fall so that I can use it in my course.