Chapter 6
Introducing Multiagent Systems to Undergraduates through Games and Chocolate

Emma Bowring
University of the Pacific, USA

Milind Tambe
University of Southern California, USA

ABSTRACT

The field of "intelligent agents and multi-agent systems" is maturing; no longer is it a special topic to be introduced to graduate students after years of training in computer science and many introductory courses in artificial intelligence. Instead, the time is ripe to introduce agents and multi-agents directly to undergraduate students, whether majoring in computer science or not. This chapter focuses on exactly this challenge, drawing on the co-authors' experience of teaching several such undergraduate courses on agents and multi-agents, over the last three years at two different universities. The chapter outlines three key issues that must be addressed. The first issue is facilitating students' intuitive understanding of fundamental concepts of multi-agent systems; the authors illustrate uses of science fiction materials and classroom games to not only provide students with the necessary intuitive understanding but with the excitement and motivation for studying multi-agent systems. The second is in selecting the right material — either science-fiction material or games — for providing students the necessary motivation and intuition; we outline several criteria that have been useful in selecting such material. The third issue is in educating students about the fundamental philosophical, ethical and social issues surrounding agents and multi-agent systems:they outline course materials and classroom activities that allow students to obtain this "big picture" futuristic vision of our science. The authors conclude with feedback received, lessons learned and impact on both the computer science students and non computer-science students.

DOI: 10.4018/978-1-60960-080-8.ch006

INTRODUCTION

Since the first international conference on multi-agent systems, ICMAS, held in 1995, to the international conference on agents and multi-agent systems (AAMAS), 2009, the entire field of "intelligent agents and multi-agent systems" as represented by AAMAS has matured significantly. In earlier years, students were introduced to this field as a special topic, only as a graduate course, after years of training in computer science, and after introductory courses in artificial intelligence. At the time, even the foundational principles of our field were unclear, and thus the only available syllabus was a set of advanced papers in multi-agent systems.

Over the years, our field has gradually matured, and we are in a historic transition period. This is similar to fields such as robotics and software engineering that have matured to a point where undergraduates in computer science are able to take courses to develop relevant skill-sets in these fields. Agents and multi-agent systems have similarly reached that critical mass. In fact, given the potential social impact of our field — society in general will need to interact with agents and multiagents on an increasing basis — introducing the fundamentals of our field to non-computer science (and non-engineering students in general) is also important.

This chapter focuses on this challenge, outlining three key issues that must be addressed. The first issue is facilitating students' intuitive understanding of fundamental concepts of multi-agent systems. There are two obstacles: (i) while a number of textbooks have been written on agents and multi-agent systems, there is not yet a standard set of concepts that are recognized as the foundation of the field and (ii) we need conceptual tools to help students understand multi-agent systems without relying on an extensive computing background. In this chapter, we outline our approach to the lack of a standardized curriculum and more importantly, to address the issue of developing intuitive understanding of our science, we introduced science fiction materials and classroom games. These tools — science fiction and games — allow us provide students with excitement about multi-agents, instilling a sense of wonder. The second issue is in selecting the right material, either science-fiction material or games: we outline several criteria that have been useful in selecting such material. In essence, we need to trade off conciseness of the material and its usefulness in the concepts taught. The key here is to ensure that a science fiction episode/story or game acts as "spice" to the main dish of the actual content from the field of agents and multi-agent systems. Thus, these extra materials should not dominate our lectures, and should help rather than hinder teaching of the desired concepts. The third issue is in educating students about the philosophical and ethical issues that have surrounded agents and multi-agent systems, as well as concerns about their future social impact. This is a crucial issue as it provides students with this bigger picture view of our field, educating them in the foundational debates such as the nature of intelligence. Science fiction is particularly important in allowing us to address this issue, providing a rich framework to construct exercises and discussions. For example, we outline a courtroom trial exercise based on a science-fiction TV episode, that allows students to debate the nature of intelligence, future rights of robots, and potential liabilities that agents designers may face in the future. Providing students with more creative and well-rounded courses that touched on philosophy, ethical and social concerns allowed us to attract a broader student audience that was keen on understanding the social context of computing.

The solutions outlined above have been practiced since 2006 via the coauthors' teaching multiple courses at two major universities in the United States to undergraduate students, both computer science majors and non-majors. For the computer science majors, we have taught three iterations of an upper-division course on multi-agent