
AC 2011-1669: AI & SCIFI: TEACHING WRITING, HISTORY, TECHNOLOGY, LITERATURE AND ETHICS

Rebecca A. Bates, Minnesota State University, Mankato

Rebecca A. Bates received the B.S. degree in biomedical engineering from Boston University in 1990, the M.S. degree in electrical engineering from Boston University in 1996 and the Ph.D. degree in electrical engineering from the University of Washington in 2004. She also received the M.T.S. degree from Harvard Divinity School in 1993. She is currently an Associate Professor in the Computer Science department at Minnesota State University, Mankato. Her research interests include speech recognition and understanding as well as engineering education.

AI & SciFi: Teaching writing, history, technology, literature and ethics

Abstract

The power of story can be used to engage students more fully in a wide range of technical topics. Stories, whether fictional or based on true case studies, have long been used to create worlds where people (or aliens) have to address ethical issues at or beyond the edges of our technical frontiers. Additionally, the history of many scientific fields coincides with their exploration in fictional works. The field of artificial intelligence (AI) has been fertile in the imaginations of authors and movie producers, with many substantive works mirroring and foreshadowing technical developments and ethical issues. This paper describes a course examining the connections between fiction and technology in the field of AI, which meets undergraduate writing, literature and ethics requirements while giving students a solid base in the AI state-of-the-art. The course can also be taught as a project-based graduate class.

Introduction

Engaging students and helping them develop creativity in engineering and science is a challenge and a goal for many faculty. Our students often come to college with extra-curricular passions that may be used to build connections with technical material. Video gaming is a common passion, but for many, the connection is to the story of the game more than to the technology. With this background, it's not surprising that many of our students connect to technology through science fiction (SciFi). This paper describes a course on artificial intelligence and science fiction. It has been offered as a general education elective, meeting writing-intensive, literature, and ethics categories for undergraduates, and as a graduate course with semester-long projects.

Artificial intelligence (AI) presents many complex theoretical, societal and ethical issues that have historically been examined in works of science fiction. The interplay between science fiction from the 1950s to the present and the development of the field of AI can be used to show how imaginative creativity and technical innovation fueled each other. As a non-AI example, the race to the moon coincided neatly with the original Star Trek television series. The interchange between creative work and technical development is especially fun to examine in the field of artificial intelligence, where there are many examples ranging from Isaac Asimov stories to *Wall-E* to intelligent artificial agents within many games.

The course goals of AI & SciFi include 1) presenting students with a fun opportunity to improve their writing, 2) exploring the social impact of the field of AI, 3) preparing students to deal with ethical questions that will arise in their professional careers, 4) connecting students to the broader culture of the technical community, 5) exploring connections between creativity in arts and in computer science, 6) developing well-spoken advocates for the field, and 7) increasing knowledge of technical content (of course). Technical topics include Turing tests, intelligent agents, artificial senses, problem solving approaches, game playing, information retrieval, machine translation and robotics. At Minnesota State University, Mankato, the course has been

approved for lower-division, general education credit in the areas of literature and ethics and also meets writing-intensive requirements because of the nature of the assignments and their required revisions. It is offered out of the computer science department rather than the English department so that there is a strong mix of technical content with ethical themes. The only other science fiction course available at the university is offered as a graduate elective for English and creative writing students.

This paper presents materials for the AI & SciFi course, both fiction and non-fiction, provides examples of assignments, describes sample student projects and presents student feedback from undergraduate and graduate offerings as well as suggestions for implementing this course.

Background

The idea of incorporating fiction in technical classes is not a new but it is one worth implementing more widely. VanderLeest and Segall both presented some of the benefits of using science fiction for engineering and technology education. Along with the built-in connections to the material that students bring to the class, important questions can be addressed related to technology's benefits and harms and who benefits or is harmed.¹ Similarly, science fiction can be used for teaching mechanics and heat transfer as well as raising issues about the ecological impact of "alien" resource use across the galaxy.² Issues related to balancing the teaching of science fiction and technology in a single course and the benefit of working with an interdisciplinary team (i.e., faculty from science and humanities) are presented by Layton, who is a member of an English department, and calls for the development of more classes like the one described here.³ He also presents issues that may make it easier for a course like this one, being taught out of a computer science department, to be approved by curriculum committees, i.e., that English departments may not be as territorial about a genre that may be considered an interloper to the field of literature.

The field of AI, and thus the fiction that incorporates ideas from AI, is rife with ethical issues (e.g., those described in Birmingham⁴). It is also, by application, an interdisciplinary field and has been used in K-12 education with philosophy of the mind as a conceptual framework to help students develop meaningful connections among AI, philosophy, science, mathematics and other fields.⁵ The historical relationship between AI fiction and technical development is illustrated in several works. For example, Vos Post & Kroeker examine the history of the development of computing technology and the relationship between fictional speculation, where authors have "tested countless technologies in the virtual environments of their fiction."⁶ Researchers recently applied the first two of Asimov's Three Laws of Robotics* to address problems in multiagent systems, particularly in the development of mixed teams of artificially intelligent agents and humans.⁷ If ethical thinking and decision making is something humans do, AI agents should have this capacity as well. Fiction (e.g., VIKI in *I, Robot* and HAL in *2001*) and ethical scholarship⁸

* The Three Laws of Robotics (Isaac Asimov): 1) A robot may not injure a human being or, through inaction, allow a human being to come to harm.

2) A robot must obey any orders given to it by human beings, except where such orders would conflict with the First Law.

3) A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

have imagined what type of agent this could be before we even have the technical capacity to create it.

AI is not the only computing field that is well-illustrated in film and novels: an overview of human-computer interaction in film is presented by Schmitz et al.⁹ Nor is AI the only field that benefits from the loop of creative imagination and technical development. Pervasive computing research has also been driven by examples like the communicators in *Star Trek*.¹⁰ (A colleague has declared for years that the point of engineering is to make life more like *Star Trek*.) An example of a similar course with a broader computing approach, including knowledge bases, web interaction and research, telepresence, virtual reality, and security as well as AI, is given by Sanderson.¹¹

Course Structure

The structure of the course allows for students to grapple with ethical questions through discussions, reflection papers and longer writing assignments. Technical content that lays the foundation for an understanding of the state-of-the-art is presented through lectures similar to those used in typical AI courses, but is tied to the implementations seen in fictional works. Works of fiction are presented in the context of foundational AI theory and state-of-the-art technical achievements. At the beginning, however, the class looks at issues of intelligence, sentience and ethical frameworks. Technically, the focus is on how to make decisions rather than a particular problem environment.

Topics covered in the course are:

- 1) Introduction to intelligence
- 2) Sentience and historical views of computing
- 3) Decision making processes and performance metrics
- 4) Agents, environments and AI vocabulary
- 5) Search algorithms
- 6) Historical comparisons, e.g., *TRON* vs. *The Computer Wore Tennis Shoes* (both produced by Disney), *TRON* vs. *WarGames* vs. *Blade Runner* (all released in 1982/3)
- 7) Programmer vs. agent responsibility
- 8) Machine learning and games, learning from examples
- 9) Examples of the 6 main disciplines of AI: natural language processing, knowledge representation, automated reasoning, machine learning, computer vision, robotics

As a four-credit semester-long undergraduate course, approximately two hours per week were spent covering technical topics and one hour was spent in group discussion. A two hour “lab” time was spent watching movies or performing exploratory labs related to technical topics. Students were asked to write weekly reflection papers on the movies and the current readings (whether technical or fictional). Students also wrote two major essays and submitted substantial revisions as revisions are a requirement for writing-intensive classes at the author’s university. The course concluded with an essay-based final exam.

As a graduate elective course taught at a foreign university on a semester system, only 4 in-class hours were available, divided into two two-hour blocks. The first session was used for the

week's movie and the second was used for lecture/discussion. Shorter films allowed some in-class discussion during the first session. Students also wrote weekly reflections on the movies, readings and technical topics and were given close feedback on their written English. The final project was structured with several milestone assignments. The project included an oral presentation, a final report examining the process of the project and the project work itself. Students were allowed to choose from three options: 1) a technical project related to AI where their final project report connected their work to the fictional works examined in class, 2) a work of fiction that incorporated a major aspect of the field of AI where the report discussed the creative process as well as technical background, and 3) an educational/outreach project that could be used to explain aspects of AI. Because all of the students in this class were working in their second (or third) language, the fictional content focused on movies, which were shown with English subtitles to reinforce understanding, rather than novels. Students were asked to read two Asimov short stories ("Runaround" and "Liar") to have a common foundation as well as one novel selected from the longer list used in the undergraduate course.

In both classes, discussion of issues and questions raised by the content was stressed. Examples of questions that fed discussion early in the course include: What is intelligence? (guided by Howard Gardner's 9 types: naturalist, musical, logical-mathematical, existential, interpersonal, bodily-kinesthetic, linguistic, intra-personal, spatial¹²) What is sentience? (having self-awareness) What is the relationship between intelligence and sentience? What are the boundaries of sentience? Who is sentient and who isn't? How do we measure intelligence? For humans? For machines?

Course Materials

The course materials consist of required and optional reading, both fiction and non-fiction, and required and optional movies and television shows. The complete list is in the Appendix. The Russell & Norvig AI text¹³ was used as a technical supplement in the graduate class and M. Tim Jones' AI text¹⁴ was used as a less dense text for the undergraduate, general education class. However, neither was required as technical content was delivered through lectures and the reading focus was on fiction. Other combinations of movies and books could also be used. For a quarter system, the number of movies and readings would need to be reduced. If there is an emphasis on AI applications and connections to fiction, *Star Wars*, *Johnny Mnemonic* and *TRON* may be more useful than *THX 1138*. If there is an emphasis on history, *Desk Set*, *WarGames*, and *I, Robot* may be more useful. For an ethics emphasis, "Measure of a Man" from *Star Trek: The Next Generation*, *Blade Runner* and *2001* may be especially useful for raising interesting questions.

Example Assignments

Weekly written reflections on the content were the primary assignments although students also had the opportunity to explore automatic decision making and machine learning through laboratory classes using the open source Orange Data Mining Toolkit.¹⁵ Assignments take into account the technical theme as well as the fictional work and addressed issues connected to historical development, ethics and societal situation. Assignments were typically presented in terms of what to notice while watching a film or reading a novel followed by questions to

address in a reflection paper. Students also had the flexibility to address issues important to them in their papers that might not be included in the assignment. Examples drawn from fiction and non-fiction course material include: How are robots treated in *Star Wars* by humans? by other robots? How is it different from treatment of humans? Are the replicants in *Blade Runner* sentient? Do the human characters think so or not? When HAL, the computer in *2001: A Space Odyssey*, kills members of the crew, who is guilty? The computer or the creator? Specific examples of assignments follow.

- 1) Watching *2001: A Space Odyssey*: This movie was released in 1968 and portrays the development of human intelligence. We see the distinction between animals and humans and then ponder the distinction between human and technological intelligence. It won an Oscar for special visual effects and was nominated for Oscars in set/art direction, best director and best screenplay. A lot of the story is told in pictures and through music. As you watch, notice in particular the interactions between the humans and HAL, the computer. Recall issues of sentience and intelligence as you notice how this agent interacts with humans. How natural are the interactions?

Reflection #1: The theme for today is machine learning. In *The Sparrow* and *2001*, we see examples of developing machine learning and the implementation of machine learning, both connected to language. One theme that shows up in *2001* is machine learning. In *The Sparrow* and *2001*, we see examples of developing machine learning and the implementation of machine learning, both connected to language. In *The Sparrow*, one of the main character's technical specialty is to analyze human processes for machine implementation. Later, we will be looking at the concepts of rule-based and statistical machine learning. Reflect on data that would match the outcome of the systems and so would be needed for training models.

Reflection #2: Reflect on the ethical issues discussed in Dennett's article "Did HAL Commit Murder" and your experience watching the movie *2001*.

- 2) Watching *War Games*: It's fun to see the differences in computing and gaming technology in this movie. It was pretty cutting edge at the time. Notice especially the ideas of games and the way that the system (WOPR) learns and communicates. There are also issues related to privacy and piracy that show up in this movie.

Watching *TRON*: This movie was released in 1982 and was considered revolutionary for its digital special effects. Even so, they still had to do some by hand because they couldn't merge the actors and the digital effects. Digital effects have come a long way since then! This movie looks at issues of control, privacy, computer interaction, security, access, and supreme beings (who is our creator? "If I don't have a user, then who wrote me?") with the backdrop of computer games where battles for power and survival took place. One of the earliest things you'll see is the personification of programs. MPC has gotten 2415 times smarter than when he was first written. Does this make him 2415 times more intelligent? More sentient? More human? At one point, the program expects the user to have a plan. Does he?

WarGames (1983) and *TRON* (1982) were released about the same time. Both films have

programs that evolved from chess-playing programs. Highlight similarities and differences. In much of our readings and viewings, we see authors working with the issue of creators and created. Include your ideas and reflections on the movie(s) and content related to games and privacy.

- 3) Watching *I, Robot*: This movie came out in 2004 and borrows a title, characters (including a company) and the 3 laws of robotics from Isaac Asimov's collection of short stories from the 1940s. What examples of AI can you see in this movie? In addition, you may think about some of the following questions How does VIKI compare to the NS-5s? How do you think self-awareness connects to intelligence? to sentience? to consciousness? Does the character Susan Calvin match the one in the book?

The theme for today is robotics. Given the readings, discussion and movies, discuss the examples of AI seen in general and in robots specifically. Write up your ideas and reflections on the movie(s) and covered content outside of class

- 4) Watching *Johnny Mnemonic*: This movie came out in 1995 and is based on a short story written in the early 1980s and a screenplay by William Gibson, author of *Neuromancer*. (Gibson coined the term "cyberspace".) Knowledge as power is a key idea in this story about a dystopian society. There is a focus on the storage and representation of data, in particular how human knowledge capacity can be augmented. Notice the methods of human computer interaction, the being that is clearly an AI, and the battle for information carried by Johnny.

Given the readings, discussion and movies, discuss the effect of data and technology on the society shown in *Johnny Mnemonic*. Is this a reasonable extrapolation from our current time? How does data relate to power and class in this society?

In general, students in the course had positive reactions to the assignments. They were engaged in discussion, gave thoughtful responses, and paid attention to detail through their use of illustrative examples.

Student Essays & Projects

The undergraduate class focused on essays that the students revised based on feedback from the instructor as well as some peer feedback. The first essay topic was quite broad but required students to use fiction and non-fiction information from the class as well as outside sources, to examine one of the following themes: 1) intelligence and its value, 2) different views of intelligence in different genres, eras or social groups, 3) intelligence vs. sentience, or 4) the development of AI and its relationship to science fiction. The second essay topic focused on an ethical issue. Students were again required to connect to fictional and non-fictional sources. Students were encouraged to compare and contrast decisions and outcomes in different sources that addressed similar issues. Students were asked to identify the stakeholders in the situation and describe the values held from the viewpoints of the different stakeholders (whether or not they were explicitly presented in the material), which may differ from the student's values, and are used in the stakeholder's decision process. Rather than "solving" the issue, the goal is for the

students to come away with a more clear understanding of the choices and actions involved. By examining the process of human decision making, students will be better equipped to develop algorithms for implementing these difficult decisions in AI agent programs.

Graduate projects were divided into three categories. There were two technical projects: one performing harmonic analysis of music using support vector machines and another detecting human vs. bot clicking for online advertising. There were two educational projects. The first examined belief revision through animated sequences using cases from *I, Robot*. The second was a web tutorial for novice users on machine learning in music recommendation systems. This project ultimately formed the basis for one student's thesis work on human interfaces for music recommendation systems. Finally, seven students chose to write their own words of fiction. Students were allowed to write fiction in their native language so peer-review was used to better evaluate the creative works. Technical topics addressed included: 1) Turing Tests and an examination of intelligence, cognition and consciousness, 2) the limits of knowledge and logic: How much can we know? 3) belief revision: What happens when robots react to changes in human world views? How should the robot deal with inconsistencies? and 4) computational analysis of humor, especially puns.

Student Feedback

Student feedback was solicited anonymously through end-of-term reports and through direct requests. Most students mentioned the importance of the technical content and their appreciation for the balance between literature and technical content as well as the writing-intensive aspect of the course.

“Reading literature and watching movies for learning was fun and motivating. Discussing and writing were great ways of stabilizing ideas in my mind.”

In both offerings, the class was “active” in that students participated in discussion, brainstorming of ideas and questions, and in directing the focus on issues pertinent to their lives. For some this was unusual:

“In this system where the demand is for results, this class had the important mission to make us think about what we do.”

“Personally I don't like traditional way of classes, and I don't think it is the best way. This course is evidence that learning can happen in a different way/structure. Academically it was rich and mind-opening (the students started to see or to question things that they didn't think of before).”

Because ethical issues can be connected to emotions, it is often easier to divorce technical content from them. However, this can be made easier when there is trust in the classroom:

“I especially liked the structure of the class and the teacher-student relationship that it sets.”

“The class' structure made me rethink what is good in a class and in a teacher-student relationship.”

At the University of São Paulo, some computer science and engineering faculty expressed concern that the class was not rigorous or technical enough, although that was not the student experience in either offering.

“The aspect of the class that was really helpful to me was the push. There was a lot of really quick reading to be done. I was forced to read faster and work a little bit harder than I was used to and it was of great benefit to me in both the short and long run.”

In terms of fostering creativity in technical innovation, one student noted that “I got from the AI & SciFi classes a much greater breadth of knowledge of the literature involved in AI. Not only that but it also aided in my ability to complexly think about AI problems. This is due to the fact that in any field there is a space that exists that can be called the total possible solution space. For the most part, people who are constantly trying to explore this space get caught up in familiar territory and have trouble ‘thinking outside the box.’ However, SciFi writers make it their job to think outside the box into this realm of total possible solutions (even if the solutions are not completely plausible, as long as they do not violate laws of nature, they are possible). This crossroads between Sci Fi and AI aided me in giving me additional out-of-the-box fodder to stimulate my thinking.”

One student comment showed that this course met the goal of fostering a connection to the culture of the larger technical community:

“I benefited ... in being able to read several great books and use them in conversation with my peers (both professional and nonprofessional). Many of the books chosen are favorites of people as nerdy as I and the ones that weren't would be if they would only read them.”

Of course, there was also critical feedback, including an acknowledgment that perhaps the details of technical content are less important than high level information. Including explicit learning goals rather than implicit goals would make it easier for students to meet the objectives.

However, the juxtaposition of technical and fictional content was important:

“I especially liked the real-world stuff that went into the class. Getting a lecture on perceptrons while reading a book in which an author makes use of the concept of Neural Networks was fascinating.”

Implementation Suggestions

As a writing-intensive course, this course can be grading-intensive for the teacher. It may also be outside of the comfort zone of faculty who focus on technical topics. Correcting grammar and structure in essays is rarely familiar for engineering and computer science faculty. However, graduate students from English or philosophy could be employed to help with grading and leading of discussions. Using graduate students to help with these aspects could also allow a class to have larger enrollment limits than the 20-25 students that are reasonable in a typical writing-intensive class. Another option is to have clear rubrics for grading essays and reflection papers as was used in these offerings. Rubrics can speed up grading and well-written rubrics can help improve student performance when given to students with the assignments by making expectations clear. Other options for reducing faculty workload while supporting student writing development include peer-reviews of writing.¹⁶

Facilitating discussion and student interaction, an important aspect of a course like this, is also not typical in science and engineering classrooms. One of the best ways to get started is to begin with smaller groups (~4-5 students) who discuss a topic or address an open-ended question and

then have each group report key findings from their discussion to the whole class. This allows for more participation by more students and removes anxiety about having to make comments in front of the class—what is reported is a consensus or a variety of points rather than an individual’s opinion. During discussion, the teacher can move from group to group, listening for a little bit and then asking questions or making brief comments to help include all members of the group or to get the group back on task if needed. Students often have experience and knowledge from other classes that can add to the discussion when they feel comfortable talking. Building relationships between students and between the teacher and the students can also make it easier to have discussions that may be more heated as the term goes on. Early in the course, it may be helpful to have a discussion of values and behaviors that allow for opinions (informed, of course) to be heard and conversation to flow. The philosopher Kwame Anthony Appiah discusses the importance of assuming that there is a chance, however small, that your opinion might not be correct in maintaining respectful conversations.¹⁷ This stance can make a difference in people’s willingness to listen and participate, even while trying to convince others that their opinion is correct.

A typical issue in engineering and computer science curricula is fitting in all the courses needed for accreditation and graduation. However, addressing ethical issues and understanding how technology fits into society as a whole are still accreditation expectations. If implementing this entire course is a problem, there can still be benefits from this approach. Including engaging fiction material can spur discussion and help students expand their ideas of how technology will fit into society. An example of incorporating science fiction into traditional AI or robotics courses is showing the movie *Star Wars* (Episode IV: A New Hope) and asking students to look at the film with a new focus. Rather than watching it as an epic battle between good and evil, students can focus on the androids in the film. This can be used to address a wide variety of issues including values, societal structure, historical imagination and expectations of the field of robotics, different areas of intelligence, implementation and the field of AI. Another example is asking students in a first-year engineering course to read a science fiction novel, like *Ender’s Game* or *The Hitchhiker’s Guide to the Galaxy*, both of which can help connect students to the broader culture of the community and raise discussion about technology, creativity and ethical situations.

Summary

The AI & SciFi course has facilitated many interesting discussions, both with students and colleagues. It has given students in majors outside of science and engineering a base of technical information while giving science and engineering students the opportunity to connect their expanding technical knowledge with interesting stories as well as ethical and historical issues. Students were asked to give anonymous feedback (which may result in clips of *2001* being shown rather than a full screening) and many reported, both directly and anonymously, that this class gave them the opportunity to discuss issues they wanted to discuss during their technical classes. Some faculty feel uncomfortable discussing philosophical or ethical issues with students, perhaps because we don’t have all of the answers. Rather than giving the answers, we can give our students a framework and context to reflect on the issues and come to their own conclusions.

Acknowledgments

The author wishes to thank the AI&SciFi students at Minnesota State University, Mankato and at the University of São Paulo for their feedback and excellent participation in the class. Discussion-based classes rely on the quality of the students so the author is grateful for the engagement students showed in the classes. She also wishes to thank the Fulbright Program for their support while on sabbatical in Brazil.

References

- [1] Steven VanderLeest, "Perspectives on Technology through Science Fiction," *Proc. American Society for Engineering Education Annual Conference and Exposition*, 2000.
- [2] Albert Segall, "Science Fiction in Engineering Instruction: To Boldly Go Where No Educator Has Gone Before," *Proc. American Society for Engineering Education Annual Conference and Exposition*, 2002.
- [3] David Layton, "Using the College Science Fiction Class to Teach Technology and Ethics: Themes and Methods," *Proc. American Society for Engineering Education Annual Conference and Exposition*, 2010.
- [4] William Birmingham, "Towards an Understanding of Artificial Intelligence and its Applications to Ethics," *Proc. American Society for Engineering Education Annual Conference and Exposition*, 2008.
- [5] Glenn W. Ellis, E.C. Ory, N. Bushan, "Organizing a K-12 AI Curriculum using Philosophy of the Mind," *Proc. American Society for Engineering Education Annual Conference and Exposition*, 2005.
- [6] Jonathan Vos Post, Kirk L. Kroeker, "Writing the Future: Computers in Science Fiction," *Computer* 33, 1, 2000, pp. 29-37.
- [7] Nathan Schurr, Pradeep Varakantham, Emma Bowring, Milind Tambe and Barbara Grosz, "Asimovian Multiagents: Applying Laws of Robotics to Teams of Humans and Agents," *Proc. 4th International Conference on Programming Multi-agent Systems*, 2007, pp. 41-55.
- [8] Nick Bostrom, "Ethical Issues in Advanced Artificial Intelligence," *Cognitive, Emotive and Ethical Aspects of Decision Making in Humans and in Artificial Intelligence*, Vol. 2, ed. I. Smit et al., Int. Institute of Advanced Studies in Systems Research and Cybernetics, 2003, pp. 12-17.
- [9] Michael Schmitz, Christoph Endres and Andreas Butz, "A Survey of Human-Computer Interaction Design in Science Fiction Movies," *Proc. 2nd International Conference on Intelligent Technologies for Interactive Entertainment*, 2007, 7:1-10.
- [10] Roy Want, "The Seeds of Inspiration," *IEEE Pervasive Computing* 7, No. 3, 2008, pp. 2-3.
- [11] Don Sanderson, "Using Science Fiction to Teach Computer Science," *Proc. WWW@10*, Terre Haute, IN, September 30-October 1, 2004.
- [12] Howard Gardner, (1983) *The Nine Types of Intelligence from Frames of Mind: the Theory of Multiple Intelligences*, Basic Books, accessed through <http://skyview.vansd.org/lshmidt/Projects/The%20Nine%20Types%20of%20Intelligence.htm>.
- [13] Stuart Russell & Peter Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed., Prentice Hall, 2010.
- [14] M. Tim Jones, *Artificial Intelligence: A Systems Approach*, Infinity Science Press, 2007.
- [15] Orange Data Mining Toolkit, available at <http://www.ailab.si/orange/>.
- [16] Richard M. Chisholm, "Introducing Students to Peer Review of Writing," *Journal on Writing Across the Curriculum* 3, No. 1, pp. 4-19.
- [17] Kwame Anthony Appiah, *Cosmopolitanism: Ethics in a World of Strangers*, W.W. Norton & Company, 2006.

Appendix: Course Materials

Required Reading

- Terry Bisson, "They're Made Out of Meat," *OMNI*, April 1991
- Orson Scott Card, *Ender's Game*, 1985
- Isaac Asimov, *I, Robot* (short stories), 1950, and *Caves of Steel*, 1953
- Neal Stephenson, *The Diamond Age*, 1995
- Mary Doria Russell, *The Sparrow*, 1996
- William Gibson, *Neuromancer*, 1984, or Neal Stephenson: *Snow Crash*, 1992
- Douglas Adams, *The Hitchhiker's Guide to the Galaxy*, 1979
- Lois McMaster Bujold, "Allegories of Change: The 'New' Biotech in the Eye of Science Fiction," *New Destinies*, Vol. VIII, 1989
- Howard Gardner, (1983) [The Nine Types of Intelligence](#) from *Frames of Mind: the Theory of Multiple Intelligences*, Basic Books.
- Alan M. Turing. (1950) [Computing Machinery & Intelligence](#), *Mind*, 59, 1950, 433-460.
- Rodney Brooks, (2008) "[I, Rodney Brooks, Am a Robot](#)", *IEEE Spectrum*, June 2008.
- Paulo Coelho, (2007) "[The Creative Process](#)", January 11, 2007
- P.W. Singer, (2009) "[Robots at War: The New Battlefield](#)", *The Wilson Quarterly*, Winter 2009.
- D. Dennett, (1997) "[Did HAL Commit Murder? \(Authorized Title\)](#)" [Unauthorized Title: "When Hal Kills, Who's to Blame? Computer Ethics"], in D. Stork, ed., *Hal's Legacy: 2001's Computer as Dream and Reality*, MIT Press 1997, p. 351-365.
- Association for Computing Machinery, [The ACM Code of Ethics](#), adopted October 16, 1992.
- Searle, John. R. (1980), "Minds, Brains and Programs", *Behavioral and Brain Sciences* 3(3): 417-457. Also in *Mind Design* (pp. 282-307). Cambridge, MA, USA: MIT Press.
- Michael Bowling, Johannes Furnkranz, Thore Graepel, Ron Musick, (2006) Machine learning and Games, *Machine Learning*, 63:211-215.
- Kelly Faircloth, (2010) "[Am I not a man and a brother: Four fictional legal systems respond to sentient androids](#)", io9.com, 7 May 2010.

Optional Reading

1. Isaac Asimov (almost anything)
2. Robert Silverberg, *The Tower of Glass*, 1970
3. Frank Herbert, *The Dosadi Experiment*, 1977
4. John W. Campbell, *The Mightiest Machine*, 1947
5. Philip K. Dick, *Do Androids Dream of Electric Sheep*, 1968
6. Stuart Russell & Peter Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed., 2009
7. M. Tim Jones, *Artificial Intelligence: A Systems Approach*, Infinity Press, 2008
8. Others as requested by students.

Required Movies

1. *Star Wars*, 1977
2. *The Computer Wore Tennis Shoes*, 1969
3. *Blade Runner*, 1982
4. *THX 1138*, 1971
5. *WarGames*, 1983
6. *2001: A Space Odyssey*, 1968
7. *Star Trek: Generations*, 1994
8. *TRON*, 1982
9. *I, Robot*, 2004
10. *Johnny Mnemonic*, 1995
11. Selected episodes of *Battlestar Galactica* and *Star Trek: The Next Generation* (e.g., “Measure of a Man”)
12. *The Hitchhiker’s Guide to the Galaxy*, (required in Brazil, optional in the US)

Optional Movies

1. *AI: Artificial Intelligence*, 2001
2. *Star Trek: The Motion Picture*, 1979
3. *Desk Set*, 1957
4. *Hitchhiker’s Guide to the Galaxy*, 2005
5. Others as requested by students.