Fiction as an Introduction to Computer Science Research

JUDY GOLDSMITH, University of Kentucky
NICHOLAS MATTEI, NICTA/UNSW

The undergraduate computer science curriculum is generally focused on skills and tools; most students are not exposed to much research in the field, and do not learn how to navigate the research literature. We describe how fiction reviews (and specifically science fiction) are used as a gateway to research reviews. Students learn a little about current or recent research on a topic that stirs their imagination, and learn how to search for, read critically, and compare technical papers on a topic related to their chosen science fiction book, movie, or TV show.

Categories and Subject Descriptors: K.3.2 [**Computer and Education**]: Computer and Informational Science Education—*Computer science education*

General Terms: Experimentation

Additional Key Words and Phrases: Artificial intelligence, research, science fiction

ACM Reference Format:
Goldsmith, J. and Mattei, N. 2014. Fiction as an introduction to computer science research. *ACM Trans. Comput. Educ.* 14, 1, Article 4 (March 2014), 14 pages.
DOI: http://dx.doi.org/10.1145/2576873

1. INTRODUCTION

Science fiction has inspired generations of would-be computer scientists and engineers. Some people draw direct lines from particular works to subsequent inventions: From Heinlein's Waldo, Inc. to modern automated assembly lines [Heinlein 1950]; from Neal Stephenson's Snow Crash to Second Life [Stephenson 1992]. Many computer science students are avid science fiction readers. This article addresses the question of how to harness their enthusiasm to propel such students past a state of passively fulfilling course obligations into the world of computer science research.

Too many undergraduate computer science students finish their education with no idea what computer science research is. Many of the students are in school, and the majority get well-paying jobs and do not pursue graduate study [Computing Research

---

The majority of the work on this system was completed while N. Mattei was at the University of Kentucky. A preliminary version of this paper was published at the 2nd AAAI Symposium on Educational Advances in Artificial Intelligence [Goldsmith and Mattei 2011].
This work is supported by the National Science Foundation, under EAGER grant CCF-1049360 and CCF-1215985. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. This work is also supported by the University of Kentucky through a Lutes Excellence in Teaching Award.
NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.
Authors' addresses: J. Goldsmith, Computer Science Department, University of Kentucky, Lexington, KY; email: goldsmit@cs.uky.edu; N. Mattei, Neville Roach Laboratory, NICTA, Sydney, Australia.

DOI: http://dx.doi.org/10.1145/2576873

Association 2009]. Many of these students have no motivation to explore the research literature, no clue about what constitutes research, peer review, or publication.

It is better for the field of computer science when practitioners have some connection to research in the field. Technology and knowledge are changing rapidly; practitioners must be able to follow trends and developments, be able to find more information about things they read about in Wired or Slashdot, and be able to evaluate the sources of information. It is important that our students be exposed to the research literature, and have at least a rudimentary understanding of the peer review process. Some may end up discovering a taste or passion for research. They may contribute to the research of other by offering technical challenges, or possibly funding. And they should be able to access and take advantage of the work that has already been done.

The goal of the exercise described here is for students to start from their own interests and to explore research related to those interests. The context of this discovery is a senior-level computer science elective, that is, Introduction to Artificial Intelligence, and the exercise consist of student reviews of science fiction books, movies, or other media that contain significant AI content. We have used the same exercise in graduate courses on AI, machine learning, and cognitive sciences, with similarly enthusiastic responses. This article focuses on the undergraduate experience.

The assignment we describe is not limited to AI, science fiction, or undergraduates. For example, students in cryptography or security courses could analyze espionage novels for their quasi-realistic portrayals of the state of the art. Students in image processing could critique the ubiquitous "enhancement" techniques shown on primetime television shows. We will describe the assignment through the lens of science fiction and artificial intelligence as it is how the assignment has been delivered to date.

For the last several years, I have walked in to the first artificial intelligence[1] class of the semester and asked the students' permission to not give exams. After a while, someone in the class usually recovers from the shock and asks what would replace the exams. I tell them that there are many options, and each student will have to choose one option for their midterm and a different one for the final. We discuss the options, and they vote. These options are in addition to 5 or 6 significant AI programming assignments that all the students are supposed to do over the course of the semester, making topics, such as search, SAT solvers, and inference and/or planning under uncertainty concrete for them.

The options on offer are that they may review a book, movie, or game with significant AI content; write a survey paper on some AI topic; implement an algorithm from the research literature; give an in-class talk on a research paper or papers; or write a short story with AI content based on actual AI research. Table I shows the distribution of student assignment choices for the 2009–2012 academic years. Approximately 2% of the students to whom I've offered these choices have mentioned considering taking an exam. None have actually taken an exam since these other options became available. All of the options have been chosen, but the book and movie reviews are by far the most popular choice, due in part to the fact that the full requirements are buried in a depth-3 web page, and students often don't click through. They think that this is the easy option, until I ask them to find the research papers to discuss the state of the art for the AI topic of interest in the book or movie. By then, they are usually committed to their proposed review, and follow through. This article will discuss the advantages of the reviewing option, and what we all bring to the process.

[1]The first author is the primary instructor for this course. Where appropriate, the first person is used to relay personal anecdotes and preferences. The second author has served as a sounding board, pedagogical expert, and occasional lecturer for the classes over several years.

Table I. Students' Project Choices 2009–2012

	2012 M	2012 F	2011 M	2011 F	2010 M	2010 F	2009 M	2009 F
SF Review	4	5	10	6	6	0	6	2
Survey	6	2	7	6	1	3	3	0
Creative Writing	2	2	1	4	0	0	2	2
Implementation	3	3	2	4	0	2	4	2
Software/Game Review	0	0	0	0	0	1	2	2
Presentations	0	0	2	1	1	1	0	1
Unrecorded	0	2	0	0	0	0	0	7
Unsubmitted	0	1	0	1	0	1	1	2

Note: Records of student midterm (M) and final (F) projects. Some projects were not submitted due to extra-academic factors such as illness.

2. BACKGROUND AND RELATED WORK

The use of AI as a hook for participation in CS is not new. AAAI has had a symposium on "Using AI to motivate greater participation in computer science" [Sahami 2008]. Similarly, the idea of AI inspiring science, and computer science in particular, is very popular (see, for example, Sawyer [2002], Watson [2003]). AAAI maintains an AI Topics site which maintains a large collection of information about AI and its use in society [AAAI 2011]. This site has been invaluable in helping to assemble this section.

This course is not the only one using science fiction to inspire students or to drive interest in artificial intelligence. Bates uses science fiction as motivation to talk about AI in a general education (nonmajors) course, and as an entry point for talking about ethics [Bates 2011]. Bowring and Tambe describe their use of science fiction and games as drivers of courses—offered to audiences as varied as students' parents; prefreshmen; and computer science majors—on multiagent systems. They use short stories and short video clips, TV shows, and movies, to present issues such as agents' models of other agents, risk-aversion, and even the use of emotions [Bowring and Tambe 2009].

The assignment described here is different from Bowring and Tambe's because they present the science fiction as class readings or in-class videos, rather than letting the students choose their own readings. Furthermore, they lead the students' explorations of the relevant technical issues, rather than pushing them to discover and explore the research literature. By being directive, Bowring and Tambe are able to focus attention on specific topics, such as distributed partially observable Markov decision processes, more deeply than our students usually do. However, our approach allows students a higher level of autonomy, thus allowing them to pursue their own interests in science fiction and AI, and giving them a flavor of self-driven research.

Courses in other fields use science fiction for motivation. Bowring and Tambe list two, including an intriguing case that used science fiction as a tool for teaching children to think about the future [Dils 1987]. I (Goldsmith) have taught a course, Science Fiction and Computer Ethics [Mihail et al. 2014].

Others have used games as a motivation and platform for exploring issues in AI [McGovern and Fager 2007]. McGovern and Fager used a gaming platform for assignments. Our students can also choose to investigate AI in games instead of science fiction. Some of them implement searching or pathfinding strategies to improve game-generated characters. However, the complexity of the programming challenges tend to limit the depth to which the students explore the AI literature. This is balanced by the students' enthusiasm for the projects, when they choose those projects. They are required to choose different forms of projects/papers/presentations for the midterm and final, so they are exposed to the research literature in at least one of those instances.

More generally, across CS there is interest in using writing both to expand student communication skills and to gain understanding and insight into the topics that the student is currently learning [Hoffman et al. 2006; Walker 1998]. There are many examples of courses where writing is emphasized as a matter of professional practice, since students must write to be successful in their careers [Pesante 1991]. There are two common approaches to writing in the "writing across the curriculum" literature [Bean 2011; McLeod and Soven 1992]: writing to learn, where the focus is on the content, and writing in the discipline, where the focus is on learning how to write (and thus to think) like a professional in the field. Some feel that these two views are incompatible [Hoffman et al. 2006] while others feel that they are closely related [Garvey 2010].

In other disciplines of computer science, instructors use writing in a variety of ways. In algorithms courses, in addition to standard "research article" type of papers, Weikle proposes assignments where students create an annotated bibliography about particular styles of algorithms [Weikle 2013]. The assignment encourages students to read broadly in a particularly focused topic and allows the students to engage with the formats and taxonomy of research publishing. Garvey discusses using writing in several different ways in upper-level programming languages courses, including both technical and creative writing to engage students and help them learn and broaden their skill set in technical writing [Garvey 2010]. One assignment, similar in spirit to the one we present, has students find and critique popular press articles with relevant research papers about programming languages.

In computer science, as with many high-consensus fields, there is a tendency to teach from a "facts and principles" standpoint [Colbeck 1998]. While many feel that this is a necessary evil in order to establish a firm foundation of core knowledge within students, there is no doubt that it reinforces "absolute truth" views of knowledge [Haworth and Conrad 1995]. Students may graduate with gaps in their ability to think and reason in situations where there may be more than one answer. While the day-to-day programming tasks of the majority of our graduates may be routine, there are still critical moments in their careers where they will need to communicate in writing about their work. In addition, they will need to think critically, for instance, to judge two or more seemingly similar technologies. We must properly train them to cut through marketing and PR material and recognize that one, the other, or both technologies may be appropriate for the task at hand. Without addressing and fostering our students' ability to think critically, we may forever leave them unable to judge multiple solutions in a fair and discerning way. Engaging multiplicity and other forms of critical thinking through exposure to research and writing will provide our students with examples of thinking that move beyond dualism and other didactic modes of reasoning [Davis 2009; Perry 1980]. We believe that exposure to multiplicity and critical thinking will better equip our students to engage in debates about technology, applications, and knowledge from multiple viewpoints.

3. THE ASSIGNMENT

The syllabus currently says the following.

> Book or movie reports will cover a work that uses AI as an integral part of the plot. You will describe the plot and the role played by AI; analyze the author(s)' concept of AI, and conclude by discussing the feasibility of this concept. Discussions of social ramifications are welcome but not required. You will be graded on the thoughtful analysis of AI in this work; readability, prose structure, and technical details of writing (spelling, punctuation, etc.) will be a non-negligible part of the grade [Goldsmith 2012c].

AI is taught as a senior-level elective. Most of the students who take it are CS majors who are out of the habit of writing, and until recently, none had taken a technical writing class. Left to their own devices, many would begin writing a day or two before the assignment was due. Therefore, proposal deadlines are set several weeks ahead. Helping students schedule and plan their writing activities is essential to ensuring success in their writing endeavor [Davis 2009].

> You will list the book/movie that you will be reviewing, cite its dominant AI theme(s) and outline the review [Goldsmith 2012b].

Proposals are submitted on paper or emailed. Email is preferred, to enable quick responses and dialogue. The goal of the discussion is to find the "hook," the AI challenge that interests them. If a student's proposal is too weak or there is not significant AI content, then I work with the student extensively to find an appropriate choice. While I have opinions on the suitability of many novels, the student has the opportunity to argue for their choice.

Once the student has chosen a work of fiction and an AI challenge, the next email interchange helps them to build a bibliography. This is an intensive process, both because it is different for each student and because students often bring up topics with which the instructor is not familiar. Ideally, the instructor should be able to point students to authors who write accessibly on the topic of interest, or at least, conferences and journals related to the topic. The goal for undergraduates is that they actually support their conclusions about feasible and possible technologies by investigating ongoing or recent research on their chosen subjects. Students are expected to read at least two peer-reviewed technical papers or chapters of advanced textbooks or monographs (graduate students are expected to read three to five papers). Once students commit to a topic, they are generally happy to learn more, even though they find that some papers are difficult to read. The midterm project accounts for 20% of the student grade while the final project accounts for 25%. I regularly check in with the students to help them move from a science fiction choice to an AI focus, bibliography, outline, and drafts (usually just the final draft). I meet or email with them about all of these, to help them meet the paper deadlines.

Grading is based on their ability to choose appropriate research sources, to link them to the science fiction, and to coherently and correctly present the state of the art (of their chosen research sources). By the time they get past the proposal, the bibliography, and the outline, the issues on which the final paper are graded are their understanding of what they've read, and their ability to present it.

Over the course of the semester, about five minutes per lecture, on average, is spent talking about research, writing, and technical writing. Some of the material is repeated each time the course is offered and some is student-driven.

3.1. Discussing Technical Papers

Before students can delve into the literature for a particular subject, they need guidance about what constitutes research literature. Before that, they need to know what research *is*. I begin with that question, and suggest that, first of all, research produces something new. The novelty can be a new idea; a new technique or algorithm; the combination of known methods in a novel manner; the application of known techniques to new problems; and analysis of algorithms or heuristics in terms of complexity or performance.

Academic scholarship can take many forms; we want students to become familiar with the scholarship of discovery [Boyer 1997]. Boyer defines four types of scholarship: discovery, integration, application, and teaching. The scholarship of discovery is the creation of novel artifacts such as hardware design, algorithms, analysis, and proofs.

Students are also encouraged to investigate the scholarship of application, that is, applying existing ideas within the field to new domains in order to extend the knowledge of the field. The scholarship of teaching, which involves research into effective teaching methodologies, is not covered in this assignment. The scholarship of integration is taking ideas from one field and applying them to another field or body of work; students learn that the act of reading and writing a review article is "research," specifically, the scholarship of integration.

For many students, the idea that "research" can be incremental is both startling and reassuring. This sometimes leads to a discussion of theses and dissertations, and of graduate school in general. This provides an opportunity to encourage students to learn more about graduate school. That conversation often occurs late in the semester, if and when they become excited about a current research area.

The next topic is quality control. I try to have recent instances of nonsensical news items or bad Wikipedia entries. The goal is not to denigrate Wikipedia, but to convince students that there is value in stable, peer-reviewed research presentations. I divide technical writing into several categories.

(1) Peer-reviewed journals and conferences
(2) Textbooks, handbooks, and monographs
(3) Tech reports and web pages
(4) Popular-press articles

Within the first category, discussion centers on the reviewing process and criteria and on the social structure of reviewing. This is an opportunity to talk about a professor's job, and the expectation that we will be on conference program committees and editorial boards. This further allows us to discuss researchers' responsibilities with respect to the research of others—a more engaging angle to the oft-repeated discussion of plagiarism, as well as what is often the students' first discussion on co-authorship.

Students ask about the reviewing process for textbooks, and often complain about books they've used. This is a springboard to remind students of the time our colleagues have spent on writing textbooks, and suggest that there is social value to buying books, whether paper or electronic.

The discussion of web pages begins with a claim that anyone can post anything. I provide several examples; once students themselves uncover examples of nonsensical web pages, they are more invested in the idea of peer review.

Next, we discuss how to find relevant articles. One goal of this exercise is to make students more information-literate: they should know not only how to navigate different sources of literature but also how to extract relevant information from what they find [Davis 2009]. They are introduced to backward referencing (reading the papers referenced in the current paper) and forward referencing (reading papers that cite the current paper). We start with CITESEER [Lawrence et al. 2005] and Google Scholar, but encourage students to look elsewhere as well.

Because students are so used to search engines, they tend to focus on individual papers. Once they have chosen their AI topics, we share links to relevant conferences and journals. They are encouraged to investigate multiple issues/instances of the journal or conference.

We discuss the value of citations: to give pointers to details and related work, and to give credit. This leads to a discussion of how to cite, how to quote, and how to choose between them. This leads to a discussion of opinions. Many students will read an opinion, stated factually, and will report that opinion as if it were their own. For instance, "Modern AI research has focused on reasoning under uncertainty," or some other broad generalization for which the students do not have sufficient information or experience

to judge. Students are encouraged to take ownership of their own opinions and to distinguish opinions from facts.

3.2. Discussing Writing

In the course of the semester, the students generally read four or more technical papers. If the AI topic that they choose is outside of my own interests, I cannot necessarily guide them to the best-written papers. Student complaints about poorly written papers are teachable moments: I ask them to analyze what makes the paper hard to read, and discuss the importance of clarity and organization.

Since some of the students are, themselves, writing technical papers (the survey option appeals to students who already have an interest in an AI topic, or who dislike genre fiction and presentations), I spend time discussing what makes a good technical paper.

What problem is being solved? Good papers should address a problem and make it clear both what the problem is and why it is important.
What have others done to solve the problem? Papers should include a survey of what others have done in order to legitimize the problem itself and to frame the new result within the existing literature.
What is new here? If the paper is about a new idea, method, or problem. It should be clear what the paper's new contribution is, and how that contribution was evaluated.
Why is it better than other solutions (or why not)? If a new idea or method is to be adopted by the greater community it must be made clear how the new method improves over old methods. Though just as valuable, but much more rare in the literature, methods that are *not* improvements need to be explained as well. By explaining clearly what went wrong and why the scientific community has learned the approaches do not work.

One of the most difficult aspects of teaching survey writing is to move students from the annotated bibliography format (one paragraph or section per paper read) to a more analytic framework. While instilling analytical writing is not the focus of the assignment, discussions on organizing technical materials pays off in the student reviews.

4. REVIEW TOPIC EXAMPLES

The quality of writing has varied. When I first gave this option, I did not insist on a bibliography of research sources. The students tended to pull conclusions from their imaginations, heavily influenced by other works of fiction and popular magazines such as *Wired*. At first I was afraid that their initial encounters with potentially impenetrable technical prose would scare them away from AI. However, I have not seen that happen. As the requirements have grown more rigorous, the papers have improved significantly, as has the students' satisfaction with the exercise.

4.1. The Bad

In the first few iterations of this assignment, I regularly got submissions on 2001 A Space Odyssey [Kubrik and Clarke 1968] and I, Robot (the original stories, usually) [Asimov 1950]. Students found 2001 more obscure than they expected. They tended to describe Asimov's ethical conundrums without much of a link to modern scholarship. I began to discourage students from choosing these books. Given the renewed interest in "robot ethics" (see, for instance, Anderson and Anderson [2007], Bryson [2010], Lin et al. [2012]); I might be willing to see that collection revisited.

Despite my best efforts, some students do not absorb the lessons about the quality of unrefereed websites, for instance. One of the earliest papers I received on I, Robot

insisted that positronic brains were, indeed feasible, "because Dana [sic], the robot in *Star Wars* [sic], has one." The reference was to what appeared to be an undergraduate paper at another university.

4.2. The Good

As with any undergraduate writing assignment, some instances are earnest and awkward. On the other hand, there are absolute gems. One student submitted a paper that began with a description of Sigourney Weaver manipulating a walking forklift (Aliens [Cameron 1986]) using her feet. The challenge there was to control a mechanism with many degrees of freedom, although the controller had only three degrees of freedom. The student then looked into applications of AI to understanding and applying complex, multidimensional motion in response to lower-dimensional controls. The AI comes in matching the sense inputs to a goal, and the goal to a set of motions or a control policy. Another student started with Babel-17 [Delaney 1966] and wrote about the Sapir-Worff hypothesis and natural language processing.

Many students ask for suggestions for novels or movies. I have suggested Hellspark [Kagan 1988] for the explicit discussion of intelligence, and received a paper on intelligent gaze control (for the robot whose intelligence is investigated). I suggested the movie, Minority Report, for the swarm, and received a paper that investigated the engineering and AI aspects of iris recognition, as portrayed in the movie. The student investigated processing speed in visual pattern recognition, as well as focal lengths for the then-current cameras. He concluded that the technology was not yet in place to pull sharp enough images out of video from cameras as far away as the scanners were in the movie. The bottleneck was not, he said, the software but the hardware. I expect that future papers on that topic will come to different conclusions.

5. MEASURES OF SUCCESS

A common measure for course success is the student evaluations of the course. In recent years, the evaluations for the first author's Introduction to Artificial Intelligence course have gone up. The recent comments from students have been uniformly enthusiastic, including, "Best class I took at UK [the University of Kentucky]. I really enjoyed doing projects over exams. I feel I learned much more that way."

One of the goals in using science fiction and requiring writing is that students will continue to be scholars long after the last class of the semester. On teacher course evaluations (TCEs) there is a specific question: "The course stimulated me to read farther in the area." Table II shows the average TCE score for this question before and after implementing the writing assignment described in this article. All responses to TCE questions from students are integers between 1 and 4. Table III shows the student's responses to the question "Rate the overall value of the course," both before and after the writing assignment was implemented. The results of a single factor ANOVA test before and after the writing assignment was implemented shows a statistically significant ($p < 0.01$) increase in student interest in continued reading. In addition to a continued interest in reading, there is a statistically significant increase in students' perception of the overall value of the course ($p < 0.05$). The enrollment of the courses has been roughly steady over the years with 12 to 17 students per semester.

While TCEs are not the most reliable instrument for gauging student involvement and course success, they are considered valid by many researchers [Marsh and Roche 1997; McKeachie 1997]. Systematic biases in TCEs have been hypothesized; however, the biasing factors are generally some combination of instructor, gender, grading leniency, whether the course is required, course workload, class size, and class time of day [Fleming et al. 2005; Guerin and Michler 2011]; all these factors are nearly constant across all the course offerings considered here.

Table II. Before and After: Reading

Without Assignment		With Assignment	
Semester	Mean	Semester	Mean
Fall 2001	3.00	Fall 2007	3.60
Fall 2003	2.90	Fall 2009	3.70
Fall 2004	2.80	Fall 2010	4.00
Fall 2005	2.60	Fall 2011	3.20
		Fall 2012	3.70

Note: Student responses to "The course stimulated me to read farther in the area." There is a statistically significant increase in the student ratings before and after the implementation of book reviews, $F(1, 7) = 24.63$, $p < 0.01$.

Table III. Before and After: Overall Value of the Course

Without Assignment		With Assignment	
Semester	Mean	Semester	Mean
Fall 2001	3.40	Fall 2007	3.40
Fall 2003	3.20	Fall 2009	3.80
Fall 2004	2.80	Fall 2010	3.90
Fall 2005	3.10	Fall 2011	3.60
		Fall 2012	3.70

Note: Student responses to "Overall value of the course." There is a statistically significant increase in the student ratings before and after the implementation of book reviews, $F(1, 7) = 6.474$, $p < 0.05$.

In 2012, over half the class chose reviews (at least 10 of the students who completed the course, and one who did not); in 2011, 15 out of 24 chose a book or movie review. (Records are incomplete, but that year there were at least four original short stories that included aspects of AI; in 2012, there were three.) In 2010, seven of eight students in the class chose book or movie reviews for their midterm project. There were two reviews of Blade Runner [Fancher and Peoples 1982], which looked at emotions and ethics; one of Daughters of Elysium [Slonczewski 2009], which looked at smart house projects; the Hacker and the Ants [Rucker 2003], which looked at swarms; the movie Moon [Jones and Parker 2009], which looked at ethics; and one review that did not actually discuss the book at all, but looked at intelligent prostheses and mental control of insect-based cyborgs. The eighth student gave a talk about game AI. This is a reasonably typical sample.

One student, writing about the movie Screamers [O'Bannon and Tejada-Flores 1995], wrote the following.

> My intention for the paper is twofold; To explain the fact vs. fiction behind the movie (what capabilities we are currently able to achieve) and to explain the social consequences, limitations, and potential benefits of the use of such weapons.

Many of the students go on to do a final project on a related subject (they are limited to at most one review). The student who looked at swarm papers for his midterm went on to implement an insect colony algorithm for his final. The student who spoke about game AI then wrote a survey paper on chess algorithms.

The students expressed great satisfaction with being able to choose the topics and formats of these projects/papers. Several have come to the class with particular interests, and have explored aspects of their topics via fiction, surveys, implementations, and presentations. All discovered something new. In all of the reviews, the fictional presentation of some aspect of AI provided direction and motivation for their reading.

I have not yet convinced any students to extend their papers for publication, although some papers have been good enough. However, most years, I recruit at least one student from the class for a research project, and several of those have led to publications or presentations.

5.1. Class Survey and Results

One of our main objectives in the design of the assignment was to encourage students to engage and interact with academic research in the field of artificial intelligence. We felt that the standard end-of-year form was not providing enough information on

Table IV. Changes in Various Questions (Totals)

	Before		After		Change (Increase)	
	Mean	Median	Mean	Median	Δ Mean	Δ Median
AI Research Is Interesting	4.46	5	4.57	5	0.11	0
Find CS Research	3.50	4	4.32	4	0.82	0
Familiar with the State of the Art	2.52	2	3.79	4	1.27	2
Perform Research	3.64	4	3.89	4	0.25	0

Note: Combined results of the questionnaire administered in the Fall 2011 and Fall 2012 academic years. The questions in **bold** text show a statistically significant increase in response before and after the course was completed ($\alpha = 0.01$).

our objectives, and so we decided to implement a supplemental survey of the students. In the Fall 2011 and Fall 2012 academic years we administered a survey to all the students enrolled in the course – a copy of the survey given on day one of the semester can be found in the appendix; in the after-semester survey, most questions reverted to the past tense. We asked a number of questions of the students that were open-ended as well as several questions that were rated on a 1 to 5 Likert scale. To compare the Likert scale data, we used Mann-Whitney U tests, which assumes equal variances but ordinal data, and used $\alpha = 0.01$, which is standard in many higher education surveys [de Winter and Dodou 2010].

In the Fall 2011 semester we had $n = 27$ surveys at the beginning of the semester and $n = 15$ at the end of the semester. In the Fall 2012 semester, the surveys started with $n = 24$ and concluded with $n = 13$. While this appears to imply that a high percentage of the students dropped the course during the semester, this is not the case. The surveys were given an IRB exemption, and they were anonymous, with no tracking information tied to the survey takers, so it is not possible to link individual student results pre and post completion of the course. Additionally, these numbers refer to the number of surveys collected on day one and on the last day of class. Thus, the initial surveys captured "class tourists," students who sign up for many classes and drop some before fees are finalized for the semester, and the final survey missed those who did not attend the last day or did not cooperate in the survey. In 2011, 22 students received grades for the course. In 2012, 17 students received grades for the course.

There is a strong survivorship bias in our data, and we are unable to filter out from the initial group the set of students who could not or did not want to complete surveys at the end of the course. As this was our first attempt at performing in-classroom research, we elected to run the survey anonymously in order to ease the IRB process. While we feel that the strength of the conclusions that can be drawn from our study is limited, we still find some interesting and supported generalizations.

In order to gauge how much the students were interested in the assignment, we asked several questions about their likes and dislikes with respect to science fiction media. Students were asked if they liked media that included artificial intelligence (before mean 4.42, median 5, and after mean 4.36, median 5). This high average score reinforces our assessment that students are not only interested in AI research, but also in media that depicts AI. Likewise, students believed that media can and does affect the direction that AI research takes (before mean 4.16, median 4, and after mean 4.21, median 4).

In response to the statement "I think research about AI is interesting," there was no significant difference before and after the course was complete ($U = 634.5$, $Z = 0.778$, $p = 0.448$, the median rank before the course was 38.19, while the median rank after the course was 62.81). However, the median score was a 5 both before and after, so the students in the course entered (and left) with a high level of respect for research in artificial intelligence. Additionally, there was no significant difference before and after

completion of the course for the statement "I want to do research in AI" ($U = 591$, $Z = 1.1896$, $p = 0.255$, the median rank before the course was 37.32, while the median rank after the course was 63.68). While the mean increased as a result of taking the course, we cannot say that there was a significant increase in interest in performing research in artificial intelligence.

We can report significant increases in the students' response to "I know how to find research on a given CS topic" ($U = 356$, $Z = 3.8319$, $p < 0.0001$, the median rank before the course was 32.62, while the median rank after the course was 68.38) and "I know about the state of the art of some AI problem" ($U = 247$, $Z = 4.8759$, $p < 0.0001$, the median rank before the course was 30.44, while the median rank after the course was 70.56) The results for these two assertions are specifically encouraging because they encompass the core of what we were attempting to highlight with the assignment.

The self-assessment statements are corroborated by open-ended response questions. At the beginning of the semester only about 30% of students named at least a vague academic resource (i.e., "a publication from a top university") as an example of where to find research on AI. At the end of the semester almost 70% of students could name a fairly specific academic resource (i.e., a specific trade organization, journal, or search engine like Google Scholar specifically for academic research). Additionally, students seemed to have a better gauge of their own abilities. At the beginning of the semester we had 17/51 (33%) students who claimed to be be able to find academic literature but were unable to name even a vague resource, and with 11/51 (21.5%) claiming and able to name a resource. At the end of the semester, 20/28 (54.9%) students claimed and were able to name a relatively specific academic resource.

Though not significant, we have examples of students shifting their future academic plans in response to the course: changing their projected plans from getting a job to attending grad school or specifically stating that they would like to obtain a PhD in computer science and focus on AI.

Despite some threats to the validity of our chosen metrics, we feel that the assignment is popular and has a positive impact on student learning outcomes. We are observing a higher level of self-reported and demonstrable skill before and after the completion of the assignment. In future semesters we hope to implement a more comprehensive system of evaluation, as we have recently with other courses, including asking students to author reflective blogs during the semester and gaining IRB approval for a more comprehensive survey, so that we can link precourse and postcourse learning outcomes for individual students [Cross and Steadman 1996; Mihail et al. 2014].

6. POINTERS TO RELEVANT SCIENCE FICTION

I maintain a list of science fiction novels with significant AI content. [Goldsmith 2012a]. In addition, AAAI maintains a web page that discusses the influences of science fiction on AI research, and vice versa [AAAI 2011]. Wikipedia has a very useful page, AI in Fiction [Wikipedia 2013a], which includes print media, movies, and TV shows. There is also Fictional Computers page which includes these media, plus comics, graphic novels, computer and video games, board games, and role-playing games [Wikipedia 2013b]. And there is Fictional Robots and Androids page with a significant non-American set of entries, particularly under graphic novels [Wikipedia 2013c]. There are others, including a particularly nice site, Robots in Films [Dirks 2010]. Also worth mentioning is the Cognitive Science Movie Index from Indiana University [Motz 2010], which explicitly states the cognitive science themes for each movie. (A student in the Intro to the Cognitive Sciences class brought this site to my attention, and several students in the class used it.)

7. CONCLUSIONS

Late in the semester, students are told that the purpose of the reviews is purely subversive: to change how they experience popular culture, particularly fictional (and popular media) portrayals of technology and AI. The goal is that they approach encounters with the authors' imaginations with both wonder and a desire to investigate. Even bad science fiction can inspire us to learn more about some facet of AI.

Good science fiction builds on the feasible and extrapolates to the possible. Good research does the same. What an author can imagine, perhaps our students can implement, now or later in their careers. A researcher's currency is new ideas. The students learn that ideas can come from fiction as well as the research literature. It biases the presentation to insist that students restrict their attention to work labeled as "science fiction," when any representation of AI in films or literature would be an equally good springboard into the research literature. It could be argued that any sufficiently realistic-seeming fiction about technology *is* science fiction. However, it is not necessary to insist on that label.

The freedom the students have to choose or discover an AI topic is often their first taste of self-guided research. Many of them conclude the semester eager to continue reading about one of their chosen topics. Typically, at least one student from the class does a research project during the next semester. Several have been coauthors on published papers, albeit unrelated to their class papers. However, the training in tracking down, reading, and making sense of research papers has served them well.

APPENDIX

SURVEY QUESTIONS

Questions 1–8 were open-ended. The students were then asked to select answers to questions 9–15 by selecting exactly one of: Strongly Disagree, Disagree, Neutral, Agree, and Strongly Agree.

(1) Are there any specific artificial intelligence research questions that interest you? What are they and briefly describe the state of the art in the area.
(2) What specific resources (books, magazines, websites, search engines) would you use to find information about new developments in artificial intelligence research?
(3) In any order list a few examples of science fiction that you enjoy (or if you dont like science fiction, list some genres that you do enjoy).
(4) In any order list a few of examples of works of science fiction media (or other genres) you like that use or present (possibly fictional) AI.
(5) Do you engage in any creative activities (programming games, writing short stories or fiction, updating Wiki articles, perform music, painting, etc.)? List these activities.
(6) Do any of these activities involve (possibly fictional) AI?
(7) What would you most like to learn about in this class?
(8) What are your plans after graduation (for example: get a job, get a masters degree, get a PhD, become a surf bum, etc.)?
(9) I like media (books, movies, TV, video games, etc.) that use AI or present (possibly fictional) AI.
(10) I think that media (books, movies, TV, video games) that use AI or present (possibly fictional) AI can affect the development of real AI and/or technology.
(11) I think research about AI is interesting.
(12) I know how to find research on a given CS topic.
(13) I know about the state of the art of some AI problem.
(14) I want to do research in AI.
(15) I want to write stories or games that use AI or present (possibly fictional) AI.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers who have made many important suggestions in expanding the scope and content of this article.

REFERENCES

AAAI. 2011. AI topics: Science fiction: Views of the future involving AI. http://www.aaai.org/AITopics/pmwiki/pmwiki.php/AITopics/ScienceFiction. (Last accessed 2011)

M. Anderson and S. L. Anderson. 2007. Machine ethics: Creating an ethical intelligent agent. *AI Mag. 28*, 4.

I. Asimov. 1950. *I, Robot*. Gnome Press.

R. A. Bates. 2011. AC 2011-1669: AI & sciFi: Teaching writing, history, technology, literature and ethics. In *Proceedings of the ASEE Annual Conference & Exposition*.

J. C. Bean. 2011. *Engaging Ideas: The Professor's Guide to Integrating Writing, Critical Thinking, and Active Learning in the Classroom*. Wiley.

E. Bowring and M. Tambe. 2009. Bridging the gap: Introducing agents and multiagent systems to undergraduate students. In *Proceedings of the Workshop on Educational Uses of Multi-Agent Systems (EDUMAS)*.

E. L. Boyer. 1997. *Scholarship Reconsidered: Priorities of the Professoriate*. Jossey-Bass.

J. J. Bryson. 2010. Robots should be slaves: Key social, psychological, ethical and design issues. In *Close Engagements with Artificial Companions,* Y. Wilks Ed., 63–74.

J. Cameron. 1986. *Aliens*. Director, James Cameron.

C. L. Colbeck. 1998. Merging in a seamless blend: How faculty integrate teaching and research. *J. Higher Edu. 69*, 6, 647–671.

Computing Research Association. 2008–2009. *Taulbee Survey*. (Last accessed 10/2010)

K. P. Cross and M. H. Steadman. 1996. *Classroom Research: Implementing the Scholarship of Teaching*. Jossey-Bass, San Francisco.

B. G. Davis. 2009. *Tools for Teaching* 2nd Ed. Jossey-Bass, San Francisco.

J. C. F. de Winter and D. Dodou. 2010. Five-point Likert items: T test versus Mann-Whitney-Wilcoxon. *Pract. Assess. Res. Eval. 15*, 11, 2.

S. R. Delaney. 1966. Babel-17. Ace.

L. S. Dils. 1987. Science fiction and the future, Yale-New Haven Teachers Institute. http://www.yale.edu/ynhti/curriculum/units/1987/2/87.02.04.x.html. (Last accessed 1987)

T. Dirks. 2010. Robots in film: A complete illustrated history of robots in the movies. http://www.filmsite.org/robotsinfilm.html. (Last accessed 2010)

H. Fancher and D. W. Peoples. 1982. *Blade Runner*. Director, Ridley Scott, based on a novel by Philip K. Dick.

R. A. Fleming, E. F. Bazen, and M. E. Wetzstein. 2005. Measuring the impact of externalities on The College of Agriculture. Teaching evaluations. *J. Agricultural Appl. Econ. 37*, 3, 635–645.

A. Garvey. 2010. Writing in an upper-level CS course. In *Proceedings of the 41st ACM Technical Symposium on Computer Science Education (SIGCSE'10)*. ACM, New York, 209–213f.

J. Goldsmith. 2012a. CS 463 AI in fiction. http://www.cs.uky.edu/~goldsmit/TOCE/books.html. (Last accessed 3/2014)

J. Goldsmith. 2012b. CS 463 Project proposals. http://www.cs.uky.edu/~goldsmit/TOCE/proposals.html. (Last accessed 3/2014)

J. Goldsmith. 2012c. CS 463 projects. http://www.cs.uky.edu/~goldsmit/TOCE/projects.html. (Last accessed 3/2014)

J. Goldsmith and N. Mattei. 2011. Science fiction as an introduction to ai research. In *Proceedings of the 2nd AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI)*.

J. T. Guerin and D. Michler. 2011. Analysis of undergraduate teaching evaluations in computer science. In *Proceedings of the ACM Technical Symposium on Computer Science Education (SIGCSE)*.

J. G. Haworth and C. F. Conrad. 1995. Curricular Transformations: Traditional and emerging voices in the academy. In *Revisioning Curriculum in Higher Education*, J. G. Haworth and C. F. Conrad Eds., Simon & Schuster, 191 –202.

R. A. Heinlein. 1950. *Waldo & Magic, Inc.* Doubleday Books.

M. E. Hoffman, T. Dansdill, and D. S. Herscovici. 2006. Bridging writing to learn and writing in the discipline in computer science education. In *Proceedings of the 37th ACM Technical Symposium on Computer Science Education (SIGCSE'06)*. 117–121.

D. Jones and N. Parker. 2009. *Moon*. Director, Duncan Jones.

J. Kagan. 1988. *Hellspark*. Tor.

S. Kubrik and A. C. Clarke. 1968. 2001: *A Space Odyssey*. Director, Stanley Kubrik.

S. Lawrence, L. Giles, K. Bollacker, and I. Councill. 2005. CITESEER. http://citeseer.ist.psu.edu/citeseer.html. (Last accessed 2005)

P. Lin, K. Abney, and G. A. Bekey. 2012. *Robot Ethics: The Ethical and Social Implications of Robotics.* MIT Press, Cambridge, MA.

H. W. Marsh and L. A. Roche. 1997. Making students' evaluations of teaching effectiveness effective: The critical issues of validity, bias, and utility. *Amer. Psychol. 52*, 11, 1187–1197.

A. McGovern and J. Fager. 2007. Creating significant learning experiences in introductory artificial intelligence. *ACM SIGCSE Bull. 39*, 1, 39–43.

W. J. McKeachie. 1997. Student ratings: The validity of use. *Amer. Psychol. 52*, 11, 1218–1225.

S. H. McLeod and M. Soven. 1992. *Writing Across the Curriculum*. Sage Publications.

P. Mihail, B. Rubin, and J. Goldsmith. 2014. Online discussions: Improving education in CS?. In *Proceedings of the 45th ACM Technical Symposium on Computer Science Education (SIGCSE'14)*.

B. Motz. 2010. Cognitive science movie index. https://www.indiana.edu/~cogfilms/. (Last accessed 2010)

D. O'Bannon and M. Tejada-Flores. 1995. *Screamers*. Director, Christian Duguay, based on short story "Second Variety" by Philip K. Dick.

W. G. Perry. 1980. Cognitive and ethical growth: The making of meaning. In *The Modern American College*, A. W. Chickering and Associates Eds., Josey Bass, 76–109.

L. H. Pesante. 1991. Integrating writing into computer science courses. In *Proceedings of the 22nd ACM Technical Symposium on Computer Science Education (SIGCSE'91)*. 205–209.

R. Rucker. 2003. *The Hacker and the Ants*: Version 2.0. Four Walls Eight Windows.

M. Sahami. 2008. Using AI to motivate greater participation in computer science. AAAI Spring Symposium. Tech. rep. SSS08, AAAI Spring Symposium.

R. J. Sawyer. 2002. AI and Sci-Fi: My, Oh, My! In *Proceedings of the 12th Annual Canadian Conference on Intelligent Systems*.

J. Slonczewski. 2009. *Daughter of Elysium - An Elysium Cycle Novel*. Phoenix Pick.

N. Stephenson. 1992. *Snowcrash*. Bantum Books.

H. M. Walker. 1998. Writing within the computer science curriculum. *ACM SIGCSE Bull. 30*, 2, 24–25.

I. Watson. 2003. The aims of artificial intelligence: A science fiction view. *IEEE Intell. Syst. 18*.

D. Weikle. 2013. Two concrete examples of upper-level writing assignments in an algorithms course. *J. Comput. Sci. Colleges 28*, 3, 14–20.

Wikipedia. 2013a. Artificial intelligence in fiction. http://en.wikipedia.org/wiki/Artificial_intelligence_in_fiction. (Last accessed 3/2013)

Wikipedia. 2013b. List of fictional computers. http://en.wikipedia.org/wiki/List_of_fictional_computers. (Last accessed 3/2013)

Wikipedia. 2013c. List of fictional robots and androids. http://en.wikipedia.org/wiki/List_of_fictional_robots_and_androids. (Last accessed 3/2013)

Received April 2013; revised October 2013; accepted November 2013