Int. J. Engng Ed. Vol. 25, No. 3, pp. 511–522, 2009
Printed in Great Britain.

0949-149X/91 $3.00+0.00

Artificial Intelligence in Pre-College Education: Learning within a Philosophy of the Mind Framework*

GLENN W. ELLIS, KATIE L. SILVA, TALIA EPSTEIN and NICHOLAS C. GIAMMARIA
Smith College, Northampton, MA 01063, USA. E-mail: gellis@smith.edu

In this paper we present an approach for integrating AI and philosophy of the mind in the K-12 classroom. Our rationale for this approach includes its capability to promote discourse and frame learning and the need to take a more holistic approach to engineering. We have developed a curriculum unit based upon the integrated approach that includes intended learning outcomes, a concept map and two example activities for the classroom. We present this curriculum unit together with the preconceptions of middle and high school students related to the concepts in the unit. The collected data shows high levels of interest in the possibilities of thinking machines, but only a naïve understanding of the issues. Some gender differences are also indicated. Finally we present two applications where the curriculum was used in the classroom. The first is a one-week unit in a middle school language arts classroom and the second is a full-year project-based high school course. Results from both settings showed high levels of engagement and were consistent with student achievement of intended learning outcomes.

Keywords: artificial intelligence; philosophy; concept map; preconceptions; pre-college

INTRODUCTION

ARTIFICIAL INTELLIGENCE (AI) is being used increasingly in both formal and informal pre-college education as a means to engage students in the study of engineering and technology. This increase is due to a variety of reasons including: the changing importance of AI in our daily lives; increasing emphasis on engineering and technology in educational frameworks; the development of robotics kits that make AI accessible to younger learners; and recognition of robotics as a useful pedagogical tool for active learning and engineering design. While we applaud the success of the robotics curricula that have been developed, we propose that supplementing them with framing activities that integrate AI and philosophy of the mind will result in the following:

- increased engagement and improved understanding by addressing student preconceptions regarding the possibilities and limitations of intelligent machines and by providing a conceptual framework to support their learning; and
- a broader and more contextualized understanding of AI by investigating ethical and societal issues related to the field—issues that play a key role in engineering and have been shown to attract a more diverse group of learners.

In this paper, we will present (1) an overview of AI in K-12 education, (2) the rationale for framing AI with philosophy of the mind, (3) example classroom activities that implement the rationale, (4) survey results identifying relevant student prior knowledge about AI (with implications for the classroom) and (5) example applications from several classroom settings.

* Accepted 30 September 2008.

Engineering and technology standards in K-12 US education

The 2002 report *Technically Speaking: Why All Americans Need to Know More about Technology* makes 'an urgent call' for increasing technical literacy in the United States [1]. The report notes that achieving this literacy depends on a more holistic understanding of technology as follows:

> Most people think that technology is little more than the application of science to solve practical problems . . . They are not aware that modern technology is the fruit of a complex interplay between science, engineering, politics, ethics, law, and other factors. People who operate under this misconception have a limited ability to think critically about technology—to guide the development and use of a technology to ensure that it provides the greatest benefit for the greatest number of citizens.

This report and many others call for a vast improvement in how K-12 students learn about engineering, technology and science. The International Technology Education Association (ITEA) and the International Society for Technology in Education (ISTE) are providing guidance by developing K-12 educational standards that address both technical content and the need for students to learn about the broader issues associated with engineering and technology. For example, the

ITEA includes several standards on 'Technology and Society' with benchmarks addressing ethical issues and a standard on the relationship between technology and other fields of study [2]. The ISTE includes standards stating that 'Students understand human, cultural, and societal issues related to technology and practice legal and ethical behavior' [3]. In addition to these organizations, individual states (beginning with Massachusetts [4]) are now creating their own engineering and technology standards.

AI in K-12 education

Although AI encompasses a broad range of topics, its implementation in K-12 education has focused largely on the topic of robotics. Hood and Hood point to one reason for this popularity: 'The use of LEGOs [robotics kits] increases the tactile and kinesthetic aspects of the learning experience and helps to make abstract concepts more concrete' [5]. Milingo et al. list the wide range of learning that can take place with robotics, including 'problem-solving skills, organization, teamwork, time management, oral and written communication skills (including technical writing), presentation skills, commitment and responsibility to self and others, good work ethics, good attitude, professionalism, and 'a design implementation capability'' [6]. The following examples illustrate their wide range of applications in the classroom.

Physical Science: Because of the ability to provide hands-on activities to support the learning process, the use of robotics (LEGO robotics in particular) has found many applications in the physical sciences. Some examples include: using LEGO robots to learn about data collection as part of a K-8 unit on robots exploring the Moon [7]; using robotics to learn about Newton's Laws of Motion by completing challenges in small groups [8]; using LEGO robots and sensors to explore basic cardiovascular mechanics in a high school after-school program [9]; and using LEGO robots to help non-English proficient students grasp the concepts of evolution compared with 'Intelligent Design' [10].

Computer Programming: Programming LEGO robots has been shown to be helpful in preparing students to learn standard programming languages. Varnado reported that the basic programming skills were more easily understood through robotics and that this made for an easier transition into more complex computer programming languages [11]. Wedward and Bruder outlined a six-week robotics program for secondary education students that included programming and exploring the components of how robots function [12].

Robot Design Challenges: Robot design challenges and competitions have become an increasingly popular way to introduce young people to the field of engineering. For example, in the FIRST Robotics Competition, teams of students and mentors from around the world are challenged to use a standard kit of parts to build a robot to meet a common design challenge [13]. An independent study of the FIRST competition showed a variety of positive outcomes including increased levels of pursuing engineering careers and increased community service [14]. Many teachers also bring the idea of robot challenges to their classrooms. For example, one teacher challenged his students to create a firefighting robot [15], based on Trinity College's Firefighting Home Robot Contest [16].

These examples are typical of the many applications of robotics in K-12 science, technology and engineering education. Unlike many of the curricula being used, Varnado describes how the study of AI can be easily integrated into many different subjects. He provides the following list:

- Art: Design an ad for the new robot lawn mower.
- Business: Research the cost effectiveness of robotics in industry.
- Computer Science: Investigate 'fuzzy logic.'
- Drama: Read and perform Karel Capek's 1921 play, *Rossum's Universal Robots*.
- English: Read robot short stories and write your own.
- Math: Explore binary systems and Boolean algebra.
- Physical Education: Compare human and robot movements.
- Psychology: Debate the issue of artificial intelligence.
- Science: Construct artificial muscles.
- Special Education: Construct a 'junk' robot.
- Technology: Build a robot. [11]

By including explicit associations between robots and humans and suggesting that the possibilities of AI should be debated, the integrated approach suggested by Varnado has parallels to the AI and philosophy curriculum introduced in this paper. However, Varnado does not specify an educational pathway for preparing students for a meaningful and rich debate on AI. In fact, we are aware of no instructional activities in the literature that explicitly ask students to wrestle with the definition of intelligence or the potential limits to artificial intelligence and the capabilities of machines. In response to this need, we present a set of intended learning outcomes for an integrated AI and philosophy curriculum unit with example learning activities and a concept map to support their achievement.

RATIONALE FOR INTEGRATING AI AND PHILOSOPHY

Using philosophy to organize technical AI content can significantly change how students view the field and result in a more holistic approach to learning and problem solving. Without philosophy, it is easier for students to think that the solution to all problems is to merely write better code or devote more resources. Philosophy shifts the question from simply 'What sensor do I need to add to make this robot distinguish between the red apple and the green apple?' to 'What is seeing? How might a sensor simulate seeing?' Such an approach has the following advantages.

Promoting Discourse and Exploration: Social cognitive theory suggests that engagement often happens in a context in which students encounter the thinking of others. Our approach explores a variety of provocative questions about the possibilities and limitations of thinking machines and their relationship to the human mind—questions that are designed to generate discourse. These questions also support self-exploration at an age that the adolescent research indicates is important for growing self-identity and introspection [17]. Exploring thinking machines becomes a means for thinking about the nature of one's own thinking, awareness and relationships with peers.

Learning within a Conceptual Framework: Research shows that when learners understand ideas within the context of a conceptual framework, they learn more effectively and are more capable of applying their knowledge to new domains [18]. Through a collaboration of engineering, education and philosophy faculty, we have developed a concept map of basic principles in AI and philosophy of the mind (see Fig. 1). The map helps students organize their understanding and meaningfully connect their new knowledge to what they already have learned. It is based on the idea of personhood—which is a useful way to structure the concepts and also fits well with the study of robotics by identifying parallels between robots and humans. In this map AI is divided into three parts—input, internal mechanisms, and output—which correspond to the human components of perception, mental processing, and response. The concept map is particularly useful for helping students visualize assumptions and their implications. For example, although separating mind from body may be a useful first step in the learning process, it can be argued that the mind may require a certain kind of body to support it (similar to software needing hardware to function properly). We help students understand these debatable relationships by identifying them with dashed arrows on the map to help students ultimately take their own stance and better identify or assimilate other perspectives. A detailed discussion of the concept map and its use in the classroom is given in Ellis et al. [19].

Engineering in a Liberal Arts Context: The advantages of understanding engineering within the context of the liberal arts is being increasingly recognized by the engineering community as neces-

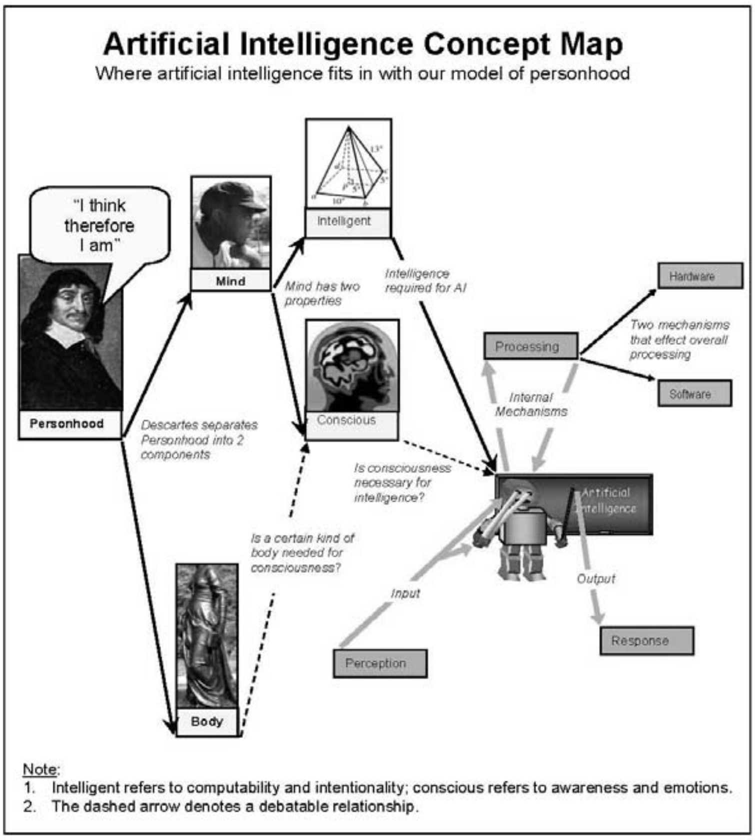

Fig. 1. Concept map organizing artificial intelligence using principles from philosophy of the mind (after Ellis et al. [19]).

sary for addressing the complex problems that engineers now face [20, 21]. This same context can also make engineering a more attractive option for girls who are underrepresented in US undergraduate engineering programs. Adelman writes that engineering would be more attractive to women if the richness of the practice—with problems brimming with ambiguities and conditional situations that are analogous to cultural context—were the framework for education [22]. Philosophy of the mind is a means for addressing questions in AI (such as the possibility of machine consciousness) that have no clear answers and lead to conditional situations that promote discussion.

Developing Metacognition Skills: Drawing parallels between humans and machines in terms of learning and knowledge is fundamental to both how engineers and scientist view AI and how philosophers view AI. These parallels present a novel opportunity for students to learn about metacognition—a key requirement for meaningful learning [18]. Learning a foreign language can help learners see their native language in new ways. Similarly, studying how machines learn, store knowledge, etc. can help students better understand their own learning. It is our intention that the study of AI be as much an exploration of human learning, intelligence, emotions and awareness as it is the study of machines.

CLASSROOM ACTIVITIES

From cognitive theory we know that learning goals are critical for helping students learn with intentionality. In Fig. 2 we specify intended learning outcomes divided into cognitions (both AI-specific and more general) and effects designed to guide the learning in an integrated AI and philosophy curriculum. To achieve these learning outcomes, we have developed a set of lesson plans and activities for the classroom in which

Cognitions—Incorporation of Ideas and Understandings

AI Specific

1. Students will understand Descartes' argument for the separation of the mind and the body (Cartesian dualism), and understand that philosophers argue about the nature of the relationship between the two.
2. Students will consider philosophy of mind and questions such as: how does a thought from the non-physical mind cause an action in the physical body, and how might computers replicate processes of the mind.
3. Students will understand the relationships inherent in the AI Concept Map. Students will consider the role of consciousness, intelligence, and life in AI.
4. Students will understand that computers must be able to meet certain criteria to be considered intelligent. These criteria may include: learning from experience and carrying on a conversation.
5. Students will understand how context shapes human language and the difficulty that AI programs have interpreting context (Searle's argument).
6. Students recognize the ethical and moral issues concerning AI.
7. Students will understand the current variety of AI applications, ranging from automatic door openers, thermostats, Internet search engines, and chatterbots.
8. Students will understand some mechanisms used by AI programs including Boolean Decision Trees and rapid decoding.
9. Students will understand that many of the issues pertaining to AI are currently being debated by scholars in the field and do not have definitive right answers.

General

1. Students will understand AI is one example of the issues currently debated by philosophers.
2. Students will understand that context is important in many situations.
3. Students will understand that moral and ethical debates often go hand-in-hand with scientific and technological developments.
4. Students will understand how ideas can be tracked across a concept map.
5. Students will understand that scientists often work in areas that do not have definitive right answers.
6. Students will understand the nature of academic discourse (following a rational argument, listening to multiple viewpoints, maintaining an objective discussion).

Effects—Behavior Goals that Reflect Certain Attitudes

1. Students will respect the opinions of others by listening attentively by engaging in objective discussion, especially regarding abstract, subjective, or ethically-charged issues.
2. Students will understand that not everyone will define consciousness, intelligence, and life in the same way. Students will expect all members of the class to define their opinions with well-formed arguments.
3. Students will gain comfort working in gray areas without a definitive 'right answer.'

Fig. 2. Intended learning outcomes for an integrated AI and philosophy of the mind unit.

students use the following pathway for developing expertise:

- Develop a definition of consciousness, intelligence and life that takes into consideration their understanding of the Concept Map.
- Develop a definition of artificial intelligence that is based upon their definition of consciousness, intelligence, and life and that incorporates what they have learned about philosophy of the mind and AI.
- Interpret the multiple and conflicting viewpoints regarding the possibilities of mind and technology.
- Make rational and justified arguments for their definitions of consciousness, intelligence, life, and the limitations for artificial intelligence that considers their experiences, the research of experts and the moral and ethical implications.
- Develop a classroom culture of academic discourse.
- Articulate their ideas in oral and written form to support their definitions, opinions, and theories about artificial intelligence and demonstrate their ability to work with abstract concepts.
- Make predictions about the future of technology and artificial intelligence.

Two example activities are described below.

AI writing prompts

The first activity is a set of writing prompts that relate to the various content areas in the AI and philosophy unit (see Fig. 3). Each of the prompts breaks down the content into a series of questions. These questions are designed to be accessible and engaging while scaffolding student responses and generating discourse. We suggest that these prompts are best utilized by providing students time to formulate their own responses before engaging in group discussions. For example, one approach would be to choose a prompt relating to the day's lesson and post it on the board so that students can begin writing immediately as they come to class. The writing time could then be followed immediately by small group or class discussion.

The prompts are adaptable to a variety of classroom situations in addition to an AI and philosophy unit. For example, they would be appropriate in a language arts class as a means of engaging students in writing about technology, engineering and philosophy—topics not typically addressed in these classes. Another possibility would be to integrate their use with a robotics curriculum or other curricula that explore issues in engineering.

Chatterbots—Interacting with AI

In Alan Turing's classic article *Computing Machinery and Intelligence* [23], he begins with the quote "I propose to consider the question, 'Can machines think?'' and then proceeds to propose a test to address the question. In the Turing Test a human interrogator engages in a text conversation with a human and a computer. The interrogator is in a different room from the subjects and only communicates with them through a text conversation (Turing originally suggested a teleprinter machine, but now there are many options available that are familiar to students). Both the human and computer try to convince the interrogator that they are human. If

1. **Defining Consciousness:**
   Define what you think makes something conscious. Do you think animals have consciousness? Do plants have consciousness? What criteria are you using to make that decision?

2. **Intelligence and Consciousness:**
   How do you define something as being intelligent? Can something be intelligent without being conscious? Explain your thinking, giving examples if possible.

3. **Consciousness and Machines:**
   Suppose you interact with someone who behaves as a normal human being does. You are able to have a coherent conversation and share jokes. Later on you are informed that this being is actually a machine. Would you consider this machine to be conscious?

4. **Creativity and Emotions:**
   Can a machine be creative? Can it have emotions? What would need to be built into the machine to do this? How would you decide whether or not the machine had accomplished this? What would it mean for your life if your computer had emotions? Would you want this to happen?

5. **Replicating the Human Mind:**
   Computers operate by using a code, or a set of rules, that determines their output in any situation. For a computer to think like a human it would have to be able to read a code that could mimic the output of every mental task, predicting a human response to every given variable. Is it possible for all of human thinking to be captured in a computer code?

6. **Technology and the Future:**
   Imagine that the year is 2200. We have the technology to recreate every living thing, down to each individual molecule. Do you think it would be possible to make a machine that could think or be conscious? Or, is there something in the mind that cannot be replicated?

Fig. 3. Writing prompts for an integrated AI and philosophy of the mind unit.

Questions

1. Rank each of the chatterbots in order of its ability to fool you into thinking that it is human. Record your results in the table below. Also, record at least three of the strengths and weaknesses of each program and your thoughts on the strategy used by each. Be sure to give each program a chance by simple questions as well as challenging it with unfair questions.

Chatterbot	Rank	Strengths	Weaknesses

2. What is your best question? Why do you consider this your best question?
3. Try telling the chatterbot something like 'I saw the mountains flying over Colorado.' Ask it if you were flying or the mountains were flying. How do they respond? Why might this sentence be challenging for a computer? What would the computer need to know to understand this question?
4. What are the strategies the chatterbots use when they do not know how to respond?
5. Try asking the computer a question with words spelled phonetically (like 'fone' for 'phone') or with intentional typos. Do these mistakes affect the chatterbot? Would you expect these kinds of errors to faze a human? Explain?
6. Tell the chatterbot a joke and then tell the same joke to your friends. How do the responses differ? What does this tell you about the computer's understanding or knowledge of what the meaning of the words are?
7. Pick any chatterbot for this question. By asking questions and doing a little online research, (many of these websites provide you with some of this information) figure out how the chatterbot works (i.e. what strategies does it use to respond to you).

Optional Expansion for Homework

1. Present the strategies used by one chatterbot on a poster.
2. Study the poster of another student who used a different chatterbot. Write a few paragraphs to compare the different strategies used by the chatterbots. Explain how these strategies are like or unlike the way you converse. Which strategy do you think is more like the way a human converses? Why? How would you use this information to improve the chatterbot program?

Fig. 4. Elements of student handout. *Meet the Chatterbots* (condensed for publication).

the judge cannot reliably tell which is which, then the machine passes the test. In this activity, students interact with chatterbots—an application of AI that attempts to simulate intelligent conversation and pass the Turing Test—to learn about criteria for machine intelligence, the importance of context in human conversation and the difficulties this poses for machines (see Fig. 2: AI Specific Intended Learning Outcomes 4, 5 and 7 and General Intended Learning Outcome 2).

Students begin the activity with a reading that introduces the topic. With the mechanics of the test understood, students are now able to start addressing the 'hook question' for the activity (Wiggins and McTighe [24] describe the use of such questions for engaging learners):

> While using Instant Messenger you suspect that you may be communicating with a computer and not a human. What question would you ask to confirm your suspicion?

Students work individually to formulate their best question and then work in groups to compare and discuss their questions—thus engaging their preconceptions through discourse with their peers and teacher. At this point, questions easily answered by a computer (i.e. 'What is your favorite color?') are typical, but some students show a more advanced understanding and formulate more challenging questions (such as asking why a certain joke is funny).

Another option in this activity is playing the gender imitation game as described by Turing [23]. In this game, instead of a human and computer trying to convince the interrogator that they are human, a male and female try to convince the interrogator that they are female (or male). We have found that this activity is fun and very engaging for students, raises issues for discussion regarding gender and stereotypes, and requires that the student interrogators practice analytical and group skills as they develop interrogation strategies. A more detailed discussion of the gender imitation game and examples from the classroom are given in Ellis and Andam [25].

At this point in the activity the students have improved their interrogation skills and are ready to converse with a variety of chatterbots that are available on the Internet [26–28]. Guided by the questions presented in Fig. 4, students compare the effectiveness of each chatterbot for carrying on a conversation. They also analyze the strengths and weaknesses of each program, try to assess how it works and note its response to a variety of questioning strategies—such as using ambiguous syntax—that require context to interpret. Many chatterbots also have extensive documentation that students can explore. We have found that students are highly engaged in this activity and in many cases develop a sophisticated understanding of the workings of the chatterbots. For example, one student so thoroughly mapped the structure of the chatterbot Alice [26], that Alice's programmers (who found his analysis that he had posted on-line during a web-search) wrote to him to express their amazement.

SURVEY RESULTS OF STUDENT AI AND PHILOSOPHY PRECONCEPTIONS

Cognitive theory posits that students need to engage with content in ways that build on what they already know. It is well understood from numerous studies that prior knowledge strongly influences the integration of new information [18]. The National Research Council (NRC) notes that '. . . teachers need to pay attention to the incomplete understandings, the false beliefs, and the naïve renditions of concepts that learners bring with them to a given subject . . .If students' initial ideas and beliefs are ignored, the understandings that they develop can be very different from what the teacher intends' [18]. Despite the increasing importance of AI in pre-college education, little is known about the prior knowledge and preconceptions that students bring to the classroom. What do students think AI is? Do they view AI as beneficial or harmful to humanity? An improved understanding of the range of preconceptions that students typically bring to the classroom can help teachers more effectively address them. To begin addressing this need, we present two data sets that provide insight into the preconceptions of middle and high school students.

JFK Middle School

The first data set was collected from the JFK Middle School in Northampton, Massachusetts, in May of 2006. A total of 81 students participated, including 37 boys and 44 girls. The data were collected by analyzing the student work produced by five sections of a seventh-grade Language Arts class at the beginning of an integrated AI and philosophy unit described later in this paper. To begin the unit, the teacher asked students to draw a picture of what they think AI is and to write a one-page essay explaining their picture and their connotative definition of the field. These drawings and explanations were examined by the authors to identify patterns and gender differences in their responses.

Examples of the drawings and essays produced by the seventh-grade students have been reported in Ellis et al. [29]. Most students indicated that AI is either a robot or a computer. Students also illustrated a positive vision of AI in almost half of the cases and at a higher rate than a negative one. In their essays students often mentioned emotions and intelligence and usually indicated that AI systems cannot have emotions or feelings. In both the essays and drawings, students typically focused on the possibilities and limitations of technology and engineering—the very issues that our AI and philosophy of the mind curriculum seeks to address. Many students indicated strong beliefs (i.e., a robot can never have feelings) and interest in these issues, but their ability to explain and justify their beliefs was usually naïve and unsupported. Engaging this high level of interest and providing a pathway to replace unsupported views with reasoned ones is central to the AI and philosophy curriculum.

Perhaps the most interesting aspect of the JFK data set is the different responses of boys and girls. Statistically significant gender differences included that only girls drew female robots/AI representation ($p = 0.00$) and boys were more likely to draw a negative portrayal of AI ($p = 0.01$). By comparison, the girls focused more on the robot's function and how it served people. Boys were also much less likely to mention either intelligence or emotion ($p = 0.00$). One implication for the classroom is that girls may find AI more engaging when it is presented in the context of helping people, since this is more consistent with the views that they bring to the classroom. This is also supported by research that has shown the importance of working on problems that are socially relevant and meaningful to attract girls to math and science [30]. It may also help to explain the often-poor representation of girls in robotics events that are based on design challenges that have no social relevance.

Smith Summer Science and Engineering Program

The second data set was collected from students participating in the 2006 Smith Summer Science and Engineering Program (SSEP). A diverse group of 75 students from grades 9 through 12 were surveyed (the survey instrument is shown in Fig. 5). The open-ended questions were designed to address six facets of understanding as described by Wiggins and McTighe [24]. These include explanation, interpretation, application, perspective, empathy and self-knowledge. Detailed results from this survey are presented in Ellis et al. [29]. The major findings of the survey are:

- *Very few students demonstrated an understanding of how a robot functions*. They showed an understanding that programming is a key issue, but in most cases they did not understand what that means.
- *Students overwhelmingly mentioned computers as an example of AI other than robots*. Often, students simply wrote 'computers' with no further explanation. Only a few gave examples that showed how AI was used on a computer.
- *Most students do not believe a computer can ever match a human in thought, action, and feeling*. While some students agreed that computers might someday think and act humanlike, only a few thought they could ever feel.

- *Students were most likely to have learned about AI through the media or in school*. Through movies in particular, the entertainment industry is creating an image of AI for students.
- *Computers are clearly a major part of these students' lives and it is through them that students see AI's effect*. Many students showed an awareness that AI has an impact on computer technology and that this directly affects them.
- *Students saw AI as being both beneficial and harmful to humanity*. For example, they reported that AI will help cut down on manual labor for humans and generally increase efficiency, but they also realized that this could lead to a shortage of jobs for humans.
- *Students noted that humans and computers have similar logic and information processing abilities*. Two examples often mentioned were analysis and memory.
- *Students had strong opinions that were about evenly split regarding harming a robot*. Some students said harming a robot would be similar to harming a living being. Others said that it would not matter because robots are not humans and do not have feelings.

The responses to the SSEP survey result in a number of implications for the classroom. Even high academic achieving girls who overwhelmingly liked their math, science and technology classes (91%) and using computers (89%) were not as motivated to learn about AI (61%). This highlights the need to provide relevant, thought-provoking hooks to engage and hold their interest. Fortunately, the philosophical issues related to AI are rich with possibilities that are intriguing, related to personal growth and can be related to technologies that affect student lives (such as voice recognition). These issues also fundamentally question how AI meets society's needs—an approach that has been shown to increase the interest of girls in math and science [30]. Finally, it is interesting to note the impact of the entertainment industry in shaping preconceptions concerning the possibilities and limitations of AI. Given the power of the medium, the potential for developing powerful misconceptions about the nature of AI is great and supports the need for identifying and addressing these misconceptions in any AI curriculum.

EXAMPLE APPLICATIONS

In this section we will present two examples of using elements from our integrated AI and philosophy unit in the classroom.

Middle School language arts

The first example is a one-week unit in a seventh-grade language arts class offered at JFK Middle School in Northampton, Massachusetts. The class was offered in five sections that totaled 81 students. The teachers chose to use the integrated AI and philosophy activities to increase engagement in the class for a greater diversity of learners. Below are the day-to-day activities of the unit:

- *Day 1: Develop a general understanding of AI and terms often associated with the field.* On the first day students brainstormed a list of films where intelligent machines mimic human behavior, discussed scientific debate in popular culture, and developed a reference sheet defining basic AI and philosophy terms.
- *Day 2: Become familiar with the work of Alan Turing and explore the nature of language.* The class began with an introduction to Alan Turing, the Turing Test, and the role of semantics and syntax in processing language. The class then worked on the chatterbot activity described in this paper.

For the following questions: SD = strongly disagree; D = disagree; N = neutral; A = agree and SA = strongly agree.

1. I like using computers. SDD N A SA
2. I am good at using computers. SDD N A SA
3. I like English/language arts classes. SD D N A SA
4. I like mathematics, science and technology class. SDD N A SA
5. I generally get good grades in school. SDD N A SA
6. I would like to learn more about AI. SD D N A SA

For most of the following questions, there are no right answers. Please just tell us what you think.

7. Explain how a robot works.
8. Robots are one example of AI. Please list other examples.
9. Some people think we can make a computer that thinks, acts and feels just like a human does. What do you think?
10. Other than this survey, how have you learned about artificial intelligence?
11. In what ways does AI affect your life now? How could it in the future
12. How can AI be beneficial to people? How could it be harmful?
13. In what ways are computers and humans alike? In what ways are they different?
14. Would it be wrong to harm a robot? Explain.

Fig. 5. SSEP Survey Questions (after Ellis et al. [29]).

- *Day 3: Understand how Boolean search engines and algorithms work.* The class used the electronic game *20 Questions* (the game tries to guess what a person is thinking by asking questions) to learn about the topics of the day. The students discussed if the game was intelligent or could think, learned about algorithms, played an online version of the game and conducted a Google search that the teacher related to the *20 Questions* game.
- *Day 4: Understand how decision trees work.* Students worked in groups to learn about decision trees and their relationship to algorithms by creating and testing a decision tree to identify a mammal based upon its characteristics.
- *Day 5: Apply AI and philosophy in a persuasive essay:* On the final day of the unit students wrote a persuasive essay voicing their support or opposition to using AI to grade student compositions for state educational testing.

The unit was assessed through a short student survey following the unit and an analysis of the Day 5 essays. Highlights from the survey indicate high levels of student interest—89% reported enjoying their language arts class more than usual during the AI unit. Most students indicated a high level of engagement and 52.4% reported exploring AI beyond the classroom requirements. Written comments indicated that they were particularly engaged in learning about how people and machines think, using computers in hands-on activities and being able to teach themselves. Most students (95.3%) reported that they had learned a lot about AI and many students (87.7%) also reported that the unit helped them better understand how people think—a positive indicator about the potential of the approach for playing a role in developing metacognitive skills.

The essays supported basic content understanding for most students. For example, 87.7% of students included in their essay a discussion about how machines have difficulty comprehending language because of context and other reasons. One student wrote 'If you have ever talked to a computer, you would find some more problems with this system, such as how it sometimes doesn't understand jokes, idioms, euphemisms, metaphors, or similes . . . Another thing our class noticed was that sometimes the machines couldn't relate pronouns to their nouns.' Other concerns discussed by students included machines lacking emotions (43.9%) and creativity (10.5%) and machines' inability to understand grammar (26.3%). Many students also reflected upon the possibilities for computers to think (72%), to be intelligent (49%) or have emotions (68%).

High School artificial intelligence

The second example is a full-year, project-based high school AI course offered at St. Paul's School in Concord, NH. The course is described in more detail with examples of student projects in Ellis and Andam [25] and included the following topics:

- Quarter 1: The course began with the AI and philosophy of the mind unit described in this paper, along with many additional topics such as the history and mathematical roots of AI. The quarter concluded with a formal class debate on the possibility of machine consciousness.
- Quarter 2: The second quarter covered an in-depth unit on connectionism in which students learned how neural networks function and developed competence modeling data with feed-forward, back-propagation ANNs. Students developed competence by completing three increasingly open-ended projects including: fruit classification, modeling housing prices and a topic chosen by the student.
- Quarters 3 and 4: Students completed an in-depth independent project that led to writing a thesis paper. The process was made as realistic as possible by requiring thesis proposals, peer review of proposals and progress reports, and a final report and formal presentation. Most students chose to pursue a research project applying ANNs. However, each year several students also decided to pursue an engineering design project related to robotics or an in-depth study of a philosophy topic related to AI.

No formal assessment was made of the course; however several informal assessments suggest the success of the approach:

- Students showed content understanding by achieving an average score of 90% on a mid-year exam encompassing content from quarters 1 and 2 (see Ellis and Andam [25] for the exam questions).
- Student interest in the course was high with about 30% of the graduating class requesting to take the course.
- Student course evaluations were 100% positive. They often cited the interesting content and interdisciplinary nature of the course, the grading emphasis on projects and papers, the flexibility to explore topics in depth that interested them, and the opportunity to conduct independent research.
- The school administration highlighted the course in recruitment literature.
- Student research resulted in a refereed article in a professional journal, a presentation at a professional conference, and several regional and national awards in science contents.

DISCUSSION

> [The liberally educated person] possesses the knowledge, not only of things, but also their mutual and true relations; knowledge, not merely considered as acquirement, but as philosophy.
>
> *John Henry Cardinal Newman* [31]

Shirley Jackson, President of Rensselaer Polytechnic Institute, comments that this 'knowledge as philosophy' approach to education is needed more today than ever for 'engineers [who] create the settings for, and means of, human interaction . . .' She cites many reasons—related to better engineering, personal growth and attracting more diverse individuals to the field—and notes that a 'liberal education ultimately makes engineers more creative by expanding their minds and exercising their imaginations' [32]. This way of thinking is becoming increasingly important as educators explore ways to prepare engineers for success in a rapidly changing global economy. If such a holistic approach to learning applies to how engineers are educated, surely it should also apply to how we educate pre-college students—most of whom will not pursue a technical career, but who will need to function and make decisions in a technical world.

A key contribution of the integrated AI and philosophy approach results from developing intended learning outcomes that include a focus not only on questions of design and application, but also on questions of theoretical possibility. As the SSEP survey showed, even students highly motivated in science and engineering have only a naïve understanding of the arguments that shape the debates about the possibilities of AI and the technology that impacts their lives. Through this approach students are pushed to define and work with challenging concepts such as intelligence and consciousness—concepts that can be more easily understood through both the application of philosophy and the study of thinking machines that engineers create. In this approach the robotic mind becomes a tool that students use for self-exploration and for making a personal connection with AI and the field of engineering, and philosophy becomes the conceptual framework that helps students make sense of what they learn. And, as the JFK students reported, such an approach also changes how students think about issues related to thought, intelligence, knowledge and understanding. With 89% of the JFK students reporting that they liked the AI unit better than their normal language arts class, it is also an engaging approach that is adaptable to a variety of situations.

Finally, exploring AI through philosophy gives students access to the ethical debate surrounding AI. Clearly engineers need to be sensitive to the social consequences of their work. Shirley Jackson writes that this sensitivity 'translates into ethics and ethical questions—i.e. it, whatever 'it' is, can be done, but should it, whatever 'it' is, be done?' [32]. Through integrating AI and philosophy, students have the opportunity to tackle the same ethical questions about engineering and technology that need to be debated by members of the AI community. By exploring questions that require probing the changing nature of the relationship between humans and computers, students are provided with many opportunities to establish and refine their own views. As illustrated by the emphasis that the girls at JFK placed upon the importance of a robot's ability to serve people, including such a focus in the integrated curriculum may also increase the interest of many students.

CONCLUSIONS

We have developed intended learning outcomes for an integrated K-12 unit on AI and philosophy of the mind along with a concept map and activities to help students successfully achieve the learning outcomes. A study of AI preconceptions of JFK middle school students showed that students possessed strong beliefs and interests in the possibilities and limitations of technology and engineering, but were generally unable to support their beliefs. The same study showed that girls tended to focus more on how AI could serve society and boys tended to focus more on the field's potential for violence. The SSEP study of high school girls showed a similar interest in and naïve understanding of AI and philosophy. Both JFK and SSEP students overwhelmingly rejected the possibility of machines possessing emotions. Based upon the content of a survey and student essays following the unit, it appears that a one-week integrated AI and philosophy course was successful at engaging students and providing learning experiences that helped them develop more reasoned positions regarding the possibilities of thinking machines. One measure of engagement showed that a little more than half of the students chose to independently explore the topic in greater depth. Data from a full-year course integrating AI and philosophy also support its potential for engagement and learning.

Acknowledgements—This research has been partially supported with a generous grant from the Ford Motor Company. We would also like to thank Lauryn McCarthy for bringing AI and philosophy into her language arts classroom.

REFERENCES

1. Committee on Technological Literacy, National Academy of Engineering and National Research Council, *Technically Speaking: Why All Americans Need to Know More About Technology*, National Academies Press, Washington, D.C., (2002).
2. International Technology Education Association. Retrieved March 15, 2008, from www.iteaconnect.org (1995).
3. International Society for Technology in Education, *National Educational Technology Standards*. Retrieved March 15, 2008, from www.iste.org/AM/Template.cfm?Section=NETS (2007).

4. Massachusetts, Department of Education, *Massachusetts Science and Technology/Engineering Curriculum Framework*, Massachusetts Department of Education, Malden (2001).
5. C. S. Hood and D. J. Hood, Teaching programming and language concepts using LEGOs, *Innovation and Technology in Computer Science Education*, 2005, pp. 19–23.
6. O. Miglino, H. H. Lundh and M. Cardaci, Robotics as an educational tool, *Journal of Interactive Learning Research*, **10**(1), 1999, pp. 25–47.
7. M. Portsmore, C. Rogers, P. Lau and E. Danahy, Remote sensing and tele-robotics for elementary and middle school via the internet, *ASEE Annual Conference Proceedings*, (2003).
8. D. C. Williams, Y. Ma, L. Prejean and M. J. Ford, Acquisition of physics content knowledge and scientific inquiry skills in a robotics summer camp, *Journal of Research on Technology in Education*, **40**(2), 2007, pp. 201–216.
9. R. Hobbs, N. Perova, I. Verner and C. Rogers, Teaching basic cardio-vascular mechanics with LEGO models: a high school case study, *ASEE Annual Conference Proceedings,* (2006).
10. E. L. Whittier and M. Robinson, Teaching evolution to non-English proficient students by using LEGO robotics, *American Secondary Education*, **35**(3), 2007, pp. 19–28.
11. T. Varnado, Robotics across the curriculum, *Tech Directions*, **60**(4), 2000, pp. 22–25.
12. K. Wedeward and S. Bruder, Incorporating robotics into secondary education, *World Automation Congress: Ninth International Symposium on Robotics with Applications,* (2002).
13. US Foundation for the Inspiration and Recognition of Science and Technology, *FIRST Robotics Competition*. Retrieved April 2, 2008, from www.usfirst.org/what/frc.
14. A. Melchior, F. Cohen, T. Cutter and T. Leavitt, *More than Robots: An Evaluation of the FIRST Robotics Competition Participant and Institutional Impacts*, Brandeis University, Waltham, (2005).
15. L. Gorman, The battling 'bots of Bloomsburg High, *Tech Directions*, **66**(6), January 2007, pp. 14–17.
16. Trinity College, *Fire Fighting Home Robot Contest*. Retrieved March 30, 2008, from www.trincoll.edu/events/robot/ (2008).
17. L. Steinberg, *Adolescence*, 7th edn, McGraw-Hill, Boston, (2007).
18. J. D. Bransford, A. L. Brown and R. R. Cocking (Eds), National Research Council Committee on Developments in the Science of Learning, *How People Learn: Brain, Mind, Experience, and School*, The National Academies Press, Washington, D.C., (2000).
19. G. W. Ellis, E.C. Ory and N. Bushan, Organizing a k-12 AI curriculum using philosophy of the mind, *ASEE Annual Conference Proceedings,* (2005).
20. ABET, Inc., *Criteria for Accrediting Engineering Programs*, Engineering Accreditation Commission, Baltimore, (2007).
21. National Science Foundation, *Moving Forward to Improve Engineering Education*, National Science Board, (2007).
22. C. Adelman, *Women and Men of the Engineering Path: A Model for Analyses of Undergraduate Careers*, National Institute for Science Education, Madison, (1998).
23. A. Turing, Computing machinery and intelligence, *Mind*, **59**(236), October 1950, pp. 433–460.
24. G. Wiggins and J. McTighe, *Understanding by Design,* 2nd edn, Association for Supervision and Curriculum Development, Alexandria, (2005).
25. G. W. Ellis and B. Andam, Teaching high school students to teach machines, *ASEE Annual Conference Proceedings,* (2004).
26. A.L.I.C.E. Artificial Intelligence Foundation, *Chat with A.L.I.C.E.* Retrieved April 1, 2008, from www.alicebot.org.
27. Eliza. Retrieved April 1, 2008, from www-ai.ijs.si/eliza-cgi-bin/eliza_script.
28. The Turing Hub, *Take the Turing Test*. Retrieved April 1, 2008, from www.turinghub.com.
29. G. W. Ellis, J. Lauer, N. Nina and K. Silva, Assessing high school girls' preconceptions about artificial intelligence to improve learning, *ASEE Annual Conference Proceedings,* (2007).
30. National Academy of Sciences and National Academy of Engineering, *Engineer Girl*. Retrieved April 11, 2008, from www.engineergirl.com (2008).
31. J. H. C. Newman, *The Idea of a University*, Yale University Press, New Haven, (1996).
32. S. A. Jackson, Engineering and the Liberal Arts, *Designing the Future: A Summit on Engineering Education at Smith College,* (2001).

Glenn W. Ellis is an Associate Professor of Engineering at Smith College. Dr. Ellis teaches courses in engineering mechanics, artificial intelligence and educational methods for teaching science and engineering. He has published numerous papers on K-16 engineering education and works with various organizations on issues of educational reform. The winner of numerous teaching awards, Dr. Ellis recently received the 2007 US Professor of the Year Award for Baccalaureate Colleges from the Carnegie Foundation for the Advancement of Teaching and the Council for Advancement and Support of Education.

Katie L. Silva graduated from Smith College in 2008 with a Bachelor of Arts in English Language and Literature and Education and Child Study. She has worked with incorporating engineering in the K-12 classroom for over two years. Her future plans are to become a secondary school teacher.

Talia Epstein helped to develop the AI and philosophy curriculum unit as a part of her thesis for the Masters of Education program at Smith College. She has experience teaching and developing curriculum for the environmental education field. She currently facilitates a peer education program to prevent alcohol and drug use in Ithaca, NY.

Nicholas C. Giammaria is a teacher of a combined 5th–6th grade class at the Greenfield Center School. His classroom combines an academic and social curriculum rooted in ethical decision making. He is a graduate of Dartmouth College with degrees in Philosophy and Environmental Studies. He served as a Peace Corps volunteer in Haiti and Armenia before receiving his Masters in education from Smith College.