Organizing a K-12 AI Curriculum using Philosophy of the Mind

Glenn W. Ellis, Eleanor C. Ory, Nalini Bhushan
Smith College, Northampton, MA

ABSTRACT

High school teams competing in the FIRST Robotics Competition and students of all ages building Lego robots have become increasingly common in K-12 education. Although robotics can be an excellent means to introduce technology and engineering design into the classroom, the scope of artificial intelligence (AI) is much greater. To help students learn about these topics in a meaningful way and to see how they fit together, this paper presents a concept map that uses principles from philosophy of the mind to organize AI topics. This approach supports a deeper understanding of AI, while making philosophical issues that interest teenagers accessible through interactive explorations of machine intelligence. In addition to the concept map, we also present examples from a K-12 AI curriculum that is being designed around the concept map.

INTRODUCTION

Lego Mindstorms robotics kits have become a well-established tool for teaching artificial intelligence (AI) in the K-12 classroom[1, 2, 3]. There is also increasing interest in using robotics and AI to change *how* children learn. For example, in *Creative Projects with LEGO Mindstorms,* Erwin writes "Dr. Papert [LEGO Professor of Learning Research at the M.I.T. Media Laboratory] and his colleagues aren't just developing hardware and software; they're developing ideas, philosophies, and theories about the nature of learning and learning environments."[1] Ellis and Andam[9] present strategies for an integrative approach to teaching AI in the K-12 environment. A key finding of their research is that integrating AI and philosophy of the mind can appeal to a broader audience and result in high student interest and achievement. In this paper we build upon Ellis and Andam's work by presenting a conceptual framework that uses philosophy of the mind to organize how students learn about AI. Our approach acknowledges that intelligence is a complex issue involving philosophical concerns that must be addressed in order to develop intelligent machines.

Using philosophy to organize AI content significantly changes how students view the field of AI. For example, the use of a philosophical framework results in a more holistic approach to learning and problem solving. Without a philosophical framework, it is easier for students to think that the solution to all problems is merely better code or more resources. Philosophy shifts the question from "What sensor do I need to add make this robot distinguish between the red apple and the green apple?" to "What is seeing? How might a sensor simulate seeing? To what extent is seeing necessary to distinguish a red apple from a green apple? Would distinguishing a red apple from a green apple be a satisfactory test?" This shift requires students to think about the "big picture" and can result in a better approach to problem solving.

The philosophical framework has another impact: it helps students organize AI content in a more meaningful way. Previous research shows that when students understand facts and ideas within the context of a conceptual framework, they learn more effectively and are more capable of applying their knowledge to new domains[5,10]. As students increase their knowledge, the philosophical framework helps them to meaningfully connect their new knowledge to what they already have learned. This approach keeps students from seeing AI as an array of separate and unconnected units; at the same time they can make explicit connections between technology and the humanities.

Based upon the success of the FIRST Robotics Competition, the LEGO Mindstorms curricula[1,2,3], the AI activities reported by Ellis and Andam[9] and the work of many other teachers and researchers, it is clear that teaching elements of AI (particularly robotics) can be engaging to pre-college students. Less obvious is that integrating AI with philosophy of the mind can make the content more interesting and relevant to a broader and more diverse group of learners. Ellis and Andam[9] found that in a high school course combining AI and philosophy, it was often the philosophical content (made accessible through hands-on AI applications) that most interested students. They observed that the AI/philosophical curriculum appeared to have been the means for self-exploration at an age when students are trying to understand the essence of their own existence and identity. The research in adolescent development supports this observation. This research indicates the importance of a growing self-identity and introspection in adolescence[6,7,8]—issues that are thoroughly explored in philosophy of the mind.

Given the idea of using philosophy as a conceptual framework for AI, what form should the framework take? One pedagogical tool for helping students organize their knowledge is the use of concept maps. Concept maps are not a new phenomenon in education. Their use is based on the theory that meaningful learning is an effortful process involving the construction of relationships between the learner's existing knowledge and new knowledge. They have been used in a wide variety of ways including for assessment[11,12,13,14], as planning tools[15,16] and for problem solving[17,18]. Figure 1 illustrates a procedure reported by Ellis et al[18]. for using teacher-generated concept maps to organize ideas in engineering curricula. Based upon this approach we have developed a concept map for philosophy of the mind that is useful for structuring AI knowledge and is developmentally appropriate for the high school classroom. The map was created by an interdisciplinary team of engineers, educators and philosophers with a review and input from the Smith College Philosophy Department. In this paper we will present the map, some of the reasoning for its structure and several activities that illustrate its use in teaching AI and philosophy through an integrated approach.

A PHILOSOPHY-BASED CONCEPT MAP FOR AI

The philosophy of mind concept map that we developed for teaching AI is shown in Figure 2. As in the development of any concept map, there are a variety of possibilities for choosing, organizing and connecting the concepts in the map. These different possibilities represent both different philosophical schools of thought as well as different ways of graphically representing the knowledge. The map design must be based upon the intended learning outcomes of the curriculum, presented in a way that is developmentally appropriate for the learner, simple

enough to not intimidate the learner, and rich enough to represent the knowledge structure and allow for further deeper exploration. The concept map shown in Figure 2 is an attempt to meet these requirements. It is based on the idea of personhood—which is both a useful way to structure the concepts and developmentally appropriate for secondary school students. A fundamental structure of the map in Figure 2 is the separation of mind and body that is grounded in Descartes' famous quote "I think therefore I am"[19]. Descartes' statement acts as a postulate for existence based on thought; he begins the analysis of intellect by abstracting thought from physical existence. This privileges one's mind over one's body in the make-up of a person. "What modern cognitive science has inherited from Descartes is the view that the mind is a private place, separated (in a number of respects) from the body and the environment that lies outside the physical boundaries on an individual"[20]. Structuring the map based upon Descartes' view of mind/body separation underscores to the learner the importance of the idea of treating properties of the mind independently of the properties of matter. This distinction acknowledges the powerful and still influential historical legacy of Descartes' vision of person.

In the concept map the mind is divided based upon two characteristics: intelligence and consciousness. Intelligence describes the 'thinking' or rational portion of mind, while consciousness describes the 'self-awareness' or 'I' portion of mind. For the sake of simplification, intelligent is defined by computability or intentionality, and conscious is defined by emotions and awareness. This representation of mind provides a framework to help students understand the debates regarding the potential and limitations of intelligent machines. These debates are fundamental to the curriculum we are developing and illustrate the advantage of a holistic approach to learning: philosophy informs AI researchers about the theoretical possibilities of machine intelligence and consciousness, while AI provides the vehicle for hands-on, contextual learning of philosophy of the mind concepts. Arguments representative of each side of the debate are presented in Figure 3.

Artificial intelligence, as the name implies, is a type of intelligence. It is represented in Figure 2 by three parts—input, internal mechanisms, and output—that correspond to the human components of perception, mental processing, and response. Perception is often mimicked using input sensors or signals. Thus, sensors that imitate human perception can be thought of as a type of intelligence. Internal processing has several components including logic. Here are based the arguments regarding whether the human mechanisms are computable. This is also the location of problem-solving skills where advances lead us to believe that technology is becoming more intelligent. The output response of an artificially intelligent mechanism is often considered the most important element for determining intelligence. For example, in the Turing Test artificial responses are compared to human responses. If the artificial response is indistinguishable from human responses, then the artificial response is considered to be intelligent.

The concept map also supports the development of a more sophisticated understanding of philosophy and AI by helping students to 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 identify these debatable relationships by identifying them with dashed arrows on the map. Another example is the separation of intelligence and consciousness. It can be argued that the mind needs a certain

emotional content or consciousness of its surroundings in order to be intelligent. Some philosophers even go so far as to deny that there is any difference between consciousness and intelligence which would negate any need for separation. What is clear is that after the concepts of artificial intelligence have been presented, students need to be able to identify where disputed definitions may take place. This allows students to ultimately take their own stance and enables them to better identify or assimilate other perspectives.

USING THE MAP IN THE CLASSROOM

An important advantage of the concept map is that it provides a conceptual framework that unites a breadth of AI technologies within one curriculum. Within this framework there are endless possibilities for students to learn about the different technologies that currently exist. Basing the map upon philosophical concepts also provides a conceptual framework for students to understand and question the technological changes that will take place in the future. The questions that are raised by AI remain the same; it is the approach or technology that is dynamic. Below we present brief descriptions of sample activities that we have developed or are developing to help students explore the ideas presented in the concept map.

Perception
In one activity students build a Lego robot that can follow a light in order to compare the behavior of a robot light sensor to a human's perception of light. Is the light sensor a fair analogy to the human eye? Students explore the point of such a robot and whether it fulfills the goal. In other activities students investigate automatic doors and thermostats and discuss whether these sensors are perceptually intelligent or not. Electronic noses and pattern recognition applications are also explored in this unit.

Processing
Students explore a wide array of topics to investigate internal processes. In one activity students learn about Bayesian decision trees and their applications (such as address books). This activity helps students contextualize the philosophy content and leads them to question whether their own mind utilizes probability mechanisms when making decisions. As part of the process students design their own Bayesian Decision Tree to determine whether a fruit is an apple or an orange. In another activity students explore the mechanics behind search engines and compare them to their own mind. Classical AI games are also used to explore some of the ways artificial intelligence poses advantages or disadvantages over human processing.

Response
As presented earlier in this paper, one standard of judging whether a machine is intelligent or not is the Turing Test. Based upon the Turing Test, Ellis and Andam describe high school activities in which students investigate the response of AI technology by conversing with "chatterbots".[9] These activities are expanded and integrated with the concept map. For example, students investigate whether they consider a judgment of intelligence solely based on the machine's response to be a sufficient metric. One activity designed to help them make this judgment is based up Searle's controversial Chinese Room argument.[31]

The Chinese Room argument posits that a completely brainless device could provide a response output that is indistinguishable from an intelligent device. In this activity, the teacher plays the role of a logician while the students pretend to be native Chinese speakers. The students are given cards with Chinese symbols, their translations and a list of translations of all possible responses. The students can slip the cards to their teacher and will receive responses in return in Chinese. After a period during which written dialogue has been exchanged between students and teacher, the class is asked whether the logician demonstrates the same fluency of Chinese as native Chinese speakers. The expected response is that students will say yes, believing that the teacher has the same translation sheets that they have. To their surprise, the logician will reveal a list of logical transcriptions with English words completely absent from the sheet (see Figure 4). This prepares them to explore Searles' arguments through a progression of questions presented in Figure 5.

DISCUSSION

Utilizing a concept map as a pedagogical tool to integrate philosophy and AI fundamentally changes how students view both fields and the relationship between them. With personhood in mind, students learn to question technologists who define the end goal as creating a machine with the equivalent response as a human. Through the map they learn visually where the debatable points in achieving this goal lie and better understand the meaning of achieving the goal. A machine that responds exactly like a human could conceivably be challenged as not intelligent on the basis that it does not perceive stimuli like a human. On the other hand, one could argue that a thermostat is intelligent on the basis that it regulates itself or that it is capable of implementing a homeostatic response to a threatening stimulus in a similar fashion to the way human beings react physiologically to an environment. However, the mere fact that a thermostat may be considered a form of artificial intelligence broadens the view of what qualifies as intelligence; a thermostat is certainly not in the same genre of AI technologies as robotics and chatterbots, but its capacity to process and respond may persuade us to consider its mechanism intelligent. The concept map helps formulate a metric for intelligence. We can take the definitions of intelligence from the map and apply them to better understand the technology.

To make the curriculum as flexible as possible, the units and activities we're developing are modular and can be used in a variety of classroom settings. For example, the concept map and any of the associated units could be used with a robotics course to help students think more deeply about the robots they are designing and to meet the learning needs of a more diverse group of learners. AI and philosophy could also be integrated with other content areas in many ways. For example, teachers could use logic activities in the mathematics classroom or natural language processing activities to study language arts. This supports our approach at Smith of using outreach programs to integrate engineering across the K-12 curriculum.

Drawing upon a broad research base, the National Research Council (NRC) recently reported the importance of metacognition (involving the learners' knowledge of their own thought process) as a key to successful learning.[10] They reported that "a metacognitive approach to instruction can help students learn to take control of their own learning by defining learning goals and monitoring their progress in achieving them." Teaching machines to learn and exploring the

nature of knowledge are important components of AI and have many connected philosophical issues. Our curriculum draws on these components to help students understand and become more aware of their own learning. Studying machine learning also makes difficult ideas more accessible. For example, in the artificial neural network unit students can see, manipulate, and measure machine learning. The AI/philosophy concept map lets them organize this knowledge in a useful way and apply it to their own learning.

In *Technically Speaking*, the National Academy of Engineering (NAE) presents the importance of increasing technological literacy for all members of our society. They point out that achieving this literacy depends on a more holistic understanding of the content areas involved in technology. They write:

> 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.[32]

Consistent with the NRC's vision, the integrated approach used in our AI/philosophy curriculum helps educate students to think differently about technology and fundamentally question the possibilities of AI for meeting society's needs.

CONCLUSION

We have developed a concept map that uses philosophical concepts to organize AI technology for use in the high school classroom. The purpose of the map is to increase learning by helping students organize their knowledge in a meaningful and holistic way. We have also developed associated activities that help students learn about the concepts presented in the map.

1. Develop initial map(s)

   Decide on the scope of the concept map(s) needed. These can be course level or program level maps that include all the major ideas and their relationships, and/or they can be more focused maps depicting, for example, problem solving strategies, unit or chapter ideas, or student's prior knowledge.

2. Introduce maps to students

   Introduce maps after an initial activity in which students identify and articulate related existing knowledge.

3. Use maps

   Refer to maps whenever new ideas are introduced to point out how the new ideas are related to ideas already learned by students. Refer to maps whenever course material is reviewed in order to make explicit and emphasize the ways the reviewed material relates to the overall course structure. Refer to maps when analyzing phenomenon of interest to show how the ideas provide a "template" or frame of reference for thinking about the phenomenon. When teaching or reviewing problems and their solutions, refer to maps in order to focus on and include strategic knowledge in classroom discourse.

4. Revise maps

   Initial maps are necessarily approximations. By engaging students with the beginning maps, they become familiar with the concept map as a tool for thought and they become participants in reshaping and refining the map to better serve their growing understanding. Refinement often adds detail, but can also result in a "master" map that is lean and shows the major relationships among ideas.

5. Repeat steps 3 & 4

Figure 1: Guidelines for Using Concept Maps (after Ellis et al.[5])

Artificial Intelligence Concept Map
Where artificial intelligence fits in with our model of person hood

"I think therefore I am"

Personhood

Mind

Intelligent[1]

Mind has two properties

Intelligence required for AI

Conscious[1]

Descartes separates Personhood into 2 components

Is consciousness necessary for intelligence?[2]

Processing

Internal Mechanisms

Hardware

Two mechanisms that effect overall processing

Software

Artificial Intelligence

Is a certain kind of body needed for consciousness?

Input

Perception

Body

Output

Response

[1]In this map we use Intelligent to refer to computability and intentionality and we use Conscious to refer to awareness and emotions

[2]The dashed arrow denotes a debatable relationship

We begin with Descartes who decided that the concept of mind can exist without a body. Mind must be both intelligent and conscious. Something intelligent perceives information, processes that information, and performs a meaningful response. In the Turing Test, a machine is considered intelligent if its response can pass as a human being. Processing can be adjusted either in the hardware, referring to physical structure and chemical makeup of the mind, or software, referring to learning and processing not requiring a physical change. Some argue that a certain kind of consciousness or awareness is required for a machine to act intelligently. It is debatable whether, in practice consciousness could exist without a special kind of body.
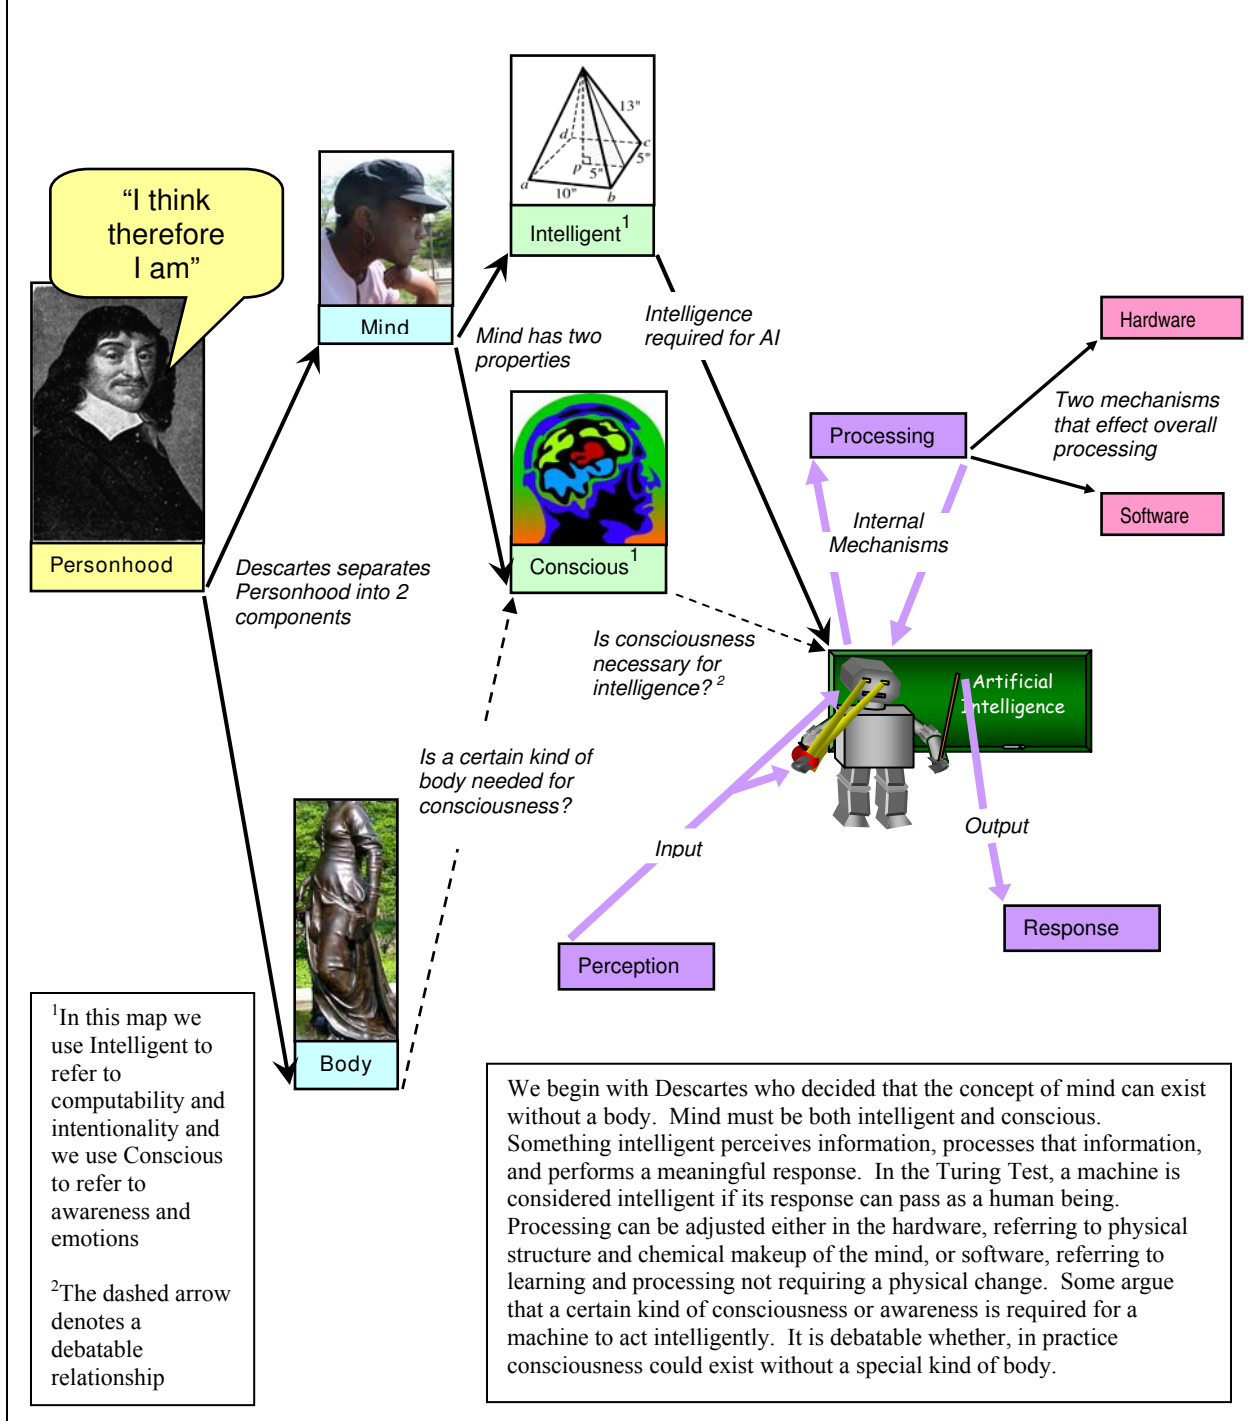
Figure 2: Concept map organizing artificial intelligence using principles from philosophy of the mind

Proponents of Strong AI
The original attempts at AI branch from the Turing machine: a machine that can solve any computable function. Proponents of strong AI believe that the mind behaves in a completely computable way:

> What does it mean to understand language? "Semantic" understanding is a correspondence between two domains; a cognitive agent understands one of those domains in terms of the other. But if one domain is to be understood in terms of another, how is the other understood? Recursively, in terms of yet another. But, since recursion needs a base case, there must be a domain that is not understood in terms of another. So, it must be understood in terms of itself. How? Syntactically! Put briefly, bluntly and a bit paradoxically, *semantic understanding is syntactic understanding.* Thus, any cognitive agent-including a computer- capable of syntax (symbol manipulation) is capable of understanding language. [21]

Strong proponents of AI whose ideas regarding the computability of intelligence are being incorporated into the curriculum include Dennett, Lycan, and Hofstadter[22,23,24,25,26,27,28].

Opponents of Strong AI
While some notable philosophers believe that intelligence can be defined as a recursively computable function, others like Searle present strong claims against this definition of intelligence. Searle begins his discussion of intentionality: "I reject any form of behaviorism or functionalism, including Turing machine functionalism, that ends up by denying the specifically mental properties of mental phenomena"[29]. In another reading, Searle goes into more detail regarding his problem with granting systems like computers intelligence:

> If I am thinking about Kansas City or wishing that I had a cold beer to drink or wondering if there will be a fall in interest rates, in each case my mental state has a certain mental content in addition to whatever formal features it might have. That is, even if my thoughts occur to me in strings of symbols, there must be more to the thought than the abstract strings, because strings by themselves can't have nay meaning. If my thoughts are to be *about* anything, then the strings must have a *meaning* which makes the thoughts about those things. In a word, the mind has more than a syntax, it has semantics.[30]

For Searle, artificial intelligence will require more than computation regardless of how similar a computed output may be to a human response.

Figure 3: Differing viewpoints on the possibilities of machine intelligence

Symbolic Manipulations (Logician)

When you receive a card from the Chinese Speaker, look for the symbols that correspond (found on the left side of the arrow). Give the Chinese Speaker the card that corresponds to the right side of the arrow

你好。 → 早晨好。

什么颜色是苹果? → 通常红色, 但有时绿色。

什么喜欢做以您的业余时间? → 读书书和在网上聊天。

什么是体育您象? → 我喜欢体育不。

我想要吃薄饼。 →
薄饼非常普遍。我喜欢意大利辣味香肠薄饼。

我发现您的反应非常乏味。 → 谢谢您的看法。

猫运行了。 → 是, 我爱他们。

我有头疼由于您。 → 我只拥有这和能只给您这。

老鼠是毛茸的。 → 关于这, 我相信例外存在。

您喜欢电影? → 我真正地喜欢科学小说和喜剧。

Figure 4: Logical manipulation sheet to be used by teachers

> Class Discussion Questions
>
> 1. How well do you think the Logician understands Chinese?
>
> 2. Now that the Logician has explained how the responses were created, think about how you converse in English. Someone tells you something, you take in this piece of information, you search in your head for an appropriate response, and you give your response to the person to whom you are conversing. In which ways are the Logician's manipulations similar to the way you converse?
>
> 3. In which ways are the Logician's manipulations unlike the way you converse? What is it about conversation that is lacking here?
>
> 4. Discuss and compare the way the Logician works and a computer works. How are they alike? How are they different?
>
> 5. Would you describe the Logician's manipulations as intelligent Chinese speaking or the Computer as an intelligent system? State very clearly what you consider to mean intelligent. (Hint: It may be helpful to look at the concept map for this question.)
>
> 6. Some people would argue that the Logician's understanding of Chinese is very similar to how a computer works. The Logician was given an input, one rule, and generated an output. Computers are said to be given an input, follow many rules step by step and generate a response. Do you think it is fair to compare the Logician with computers? Why or Why not?
>
> 7. Do you think this exercise proves that a computer can never be intelligent? Why or why not?

Figure 5: Questions for class discussion, given to students after teacher shows student the symbolic manipulation sheet

REFERENCES

[1]B. Erwin, *Creative Projects with LEGO Mindstorms*, Addison-Wesley, Boston (1998).

[2]D. Baum, *Dave Baum's Definitive Guide to LEGO Mindstorms,* APress, Emeryville (2000).

[3]M. Meadhra, P.J. Stouffer, *LEGO Mindstorms for Dummies,* IDG Books Worldwide, New York (2001).

[4]D. Grasso, "Engineering a Liberal Education,"Prism, Nov. (2002).

[5] G.W. Ellis, Al Rudnitsky, Becky Silverstein, "Using conceptual Maps to Enhance Understanding in Engineering Education," 2004 *American Society for Engineering Education Annual Conference & Exposition*, Salt Lake City Utah (2004).

[6] K. Woodward, *Understanding Identity*, Arnold, London, 1-11 (2002).

[7] J.J. Arnett, *Readings on Adolescence and Emerging Adulthood*, Prentice Hall, Upper Saddle River 111-120 (2002).

[8] C. Moore, K. Lemmon, *The Self in Time*, Lawrence Erlbaum Associates, Mahwah, 15-35, 97-121 (2001).

[9] G.W. Ellis, B. Andam, "Teaching High School Students to Teach Machines" *American Society for Engineering Education Annual Conference & Exposition*, Salt Lake City Utah (2004).

[10] NRC Commission on Behavioral and Social Sciences and Education, *How People Learn*, National Academy Press, Washington, D.C. (2000).

[11] Wallace, J.D. and J.J. Mintzes, "The concept map as a research tool: Exploring conceptual change in biology," *Journal of Research in Science Teaching*, 27, 10, (1990).

[12] Walker, J.M.T. and P.H. King, "Concept mapping as a form of student assessment and instruction in the domain of bioengineering," *Journal of Engineering Education*, 92, 2, (2003).

[13] Besterfield-Sacre, M., J. Gerchak, L.J. Shuman and H. Wolfe, Using Concept Maps for Evaluating Program Objectives, Proceedings, *33rd ASEE/IEEE Frontiers in Education Conference*, Boulder, CO, Nov. 5-8, (2003).

[14] Starr, M.L. and J.S. Krajcik, "Concept maps as a heuristic for science curriculum development: Toward improvement in process and product," *Journal of Research in Science Teaching*, 27, 10, (1990).

[15] Posner, G. and A. Rudnitsky, *Course Design: A Guide to Curriculum Development for Teachers*, 6th Edition, Allyn & Bacon Longman, Boston, (2000).

[16] Pankratius, W.J., "Building an organized knowledge base: Concept mapping and achievement in secondary school physics," *Journal of Research in Science Teaching*, 27, 10, (1990).

[17] Ellis, G.W. and Turner, W., "Helping students organize and retrieve their understanding of dynamics," Proceedings, *2003 American Society for Engineering Education Annual Conference and Exposition*, Nashville, Tennessee, June 22-25, (2003).

[18] Ellis, G.W., A. Rudnitsky and B. Silverstein, "Using Concept Maps to Enhance Understanding in Engineering Education," *International Journal of Engineering Education* (in press).

[19] Descartes, R. translated by Laurence, J.L., *Meditations on First Philosophy*, Macmillan Publishing Company, New York (1989).

[20] Butler, K., "The Scope of Psychology" *Proceedings Biennial Meeting of the Philosophy of Science Association*, East Lansing, Vol 1 428-436, (1994).

[21] Rapaport, W.J., "Understanding Understanding Syntactic Semantics and Computational Cognition" *Philosophical Perspectives*, Vol. 9 49-88, (1995).

[22] Lycan, W.G., *Judgement and Justification* Cambridge University Press, New York (1988).

[23] Dennett, D.C., *The Intentional Stance* The MIT Press, Cambridge (1987).

[24] Dennett, D.C., *Consciousness Explained* Little, Brown and Company, Boston (1991).

[25] Dennett, D.C., *Brainchildren* The MIT Press, Cambridge (1998).

[26] Dennett, D.C., *Brainstorms* Bradford Books, Publishers, Cambridge (1978).

[27] Hofstadter, D.R., *Godel, Escher, Back: an Eternal Golden Braid* Basic Books, Inc., Publishers, New York (1979).

[28] Kurzweil, R., *The Age of Intelligent Machines* The MIT Press, Cambridge (1990).

[29] Searle, J.R., *Intentionality* Cambridge University Press, Cambridge (1983).

[30] Searle, J.S., *Minds, Brains and Science* Harvard University Press, Cambridge (1984).

[31] Searle, J.R. "Minds, brains, and programs" *The Behavioral and Brain Sciences, Vol. 3, 417-457 (1980).*

[32] Pearson, G. and A.T. Young, *Technically Speaking: Why All Americans Need to Know More About Technology*, National Academy Press, Washington, DC, (2002).

BIOGRAPHICAL INFORMATION

GLENN W. ELLIS is the Ford Motor Visiting Professor of Engineering Education in the Picker engineering Program at Smith College. Educated as an engineer, he taught at both the college and secondary school levels. His current research is developing K-16 engineering curricula that adhere to the best research on learning and teaching.

ELEANOR C. ORY is a junior engineering and philosophy student in the Picker Engineering Program at Smith College. She has been working in engineering education research for one year.

NALINI BHUSHAN is an Associate Professor of Philosophy at Smith College. Her interests lie at the intersection of philosophy of the mind and language.