10-25-2004

Teaching High School Students to Teach Machines

Glenn W. Ellis
Smith College, gellis@smith.edu

Baaba Andam
Smith College

Session # 3630

Teaching High School Students to Teach Machines

Glenn W. Ellis and Baaba Andam
Smith College, Northampton, MA

ABSTRACT

In this paper we present elements of a learner-centered AI curriculum for high school students that was field-tested in two private high schools. One of these elements is a unit in which students explore the possibilities of machine intelligence and consciousness through readings and hands-on activities. Also presented in the paper is a unit for teaching students about artificial neural networks (ANN) and their application. In this unit students learn to develop and train ANNs through small projects and activities that lead up to an independent research project. Examples of student projects are presented including the application of ANNs for modeling the ozone disinfection of water, the price of real estate as determined by housing features in a local market, and admission into an independent boarding school based upon admissions application data.

INTRODUCTION

The field of robotics has becoming increasingly accessible to students of all ages. For example, Lego robotics kits are now commonly used to teach robotics in K-12 classrooms. However, other branches of artificial intelligence (AI) still remain inaccessible to pre-college students in spite of their educational potential.

In this paper we will present our curriculum for two topics covered in a learner-centered AI course that we have developed. The first topic is philosophy of the mind with a focus on the possibility of machine consciousness and the Turing test. Our experience has shown that high school students are particularly interested in the questions investigated in this part of the course. AI may be the means for self-exploration at a time when students are trying to understand the essence of their own existence and identity. This topic also provides part of a conceptual framework into which the technical subjects of the course can be understood and examined. The second topic presented is artificial neural networks (ANNs) and their application for student research. Because of their ability to learn and importance in the AI field, we have found the study of ANNs to be engaging to students and to offer numerous possibilities for development of metacognitive and critical thinking skills.

The curriculum described in this paper has been taught at St. Paul's School, a private, independent boarding school in Concord, NH and Brunswick School, a private, independent day school in Greenwich, CT. At St. Paul's School the class was taught for four years with a class size after the first year of 12 students. Because of the demand for the course and the course registration procedure, almost all of the students were seniors. At Brunswick School the class was taught for three years to a class size of about 6 students. At Brunswick School students from grades 9-12 were equally represented.

COURSE CONTENT

In the fall semester students are introduced to philosophy of the mind as they explore machine consciousness. These issues are revisited and applied to content presented throughout the rest of the course. Next, the students explore the history of artificial intelligence. While they do not write computer code, they do use a variety of classical AI software applications. Then students learn about the Turing Test and apply the ideas behind it in a number of applications. Finally, students complete an in-depth unit on connectionism in which they learn about and use artificial neural networks (ANNs).

In the spring semester students apply their understanding of artificial neural network through an in-depth independent research topic. While working on their projects, students also explore AI topics such as robotics or artificial life. Curriculum elements from several topics throughout the course are presented in greater depth in the following sections.

Machine Consciousness

> There's something queer about describing consciousness: whatever people mean to say, they just can't seem to make it clear. It's not like feeling confused or ignorant. Instead, we feel we know what's going on but can't describe it properly. How could anything seem so close, yet always keep beyond our reach?
>
> --Marvin Minsky, *The Society of Mind*[1]

Students begin the course by investigating the question, "Can machines be conscious?" To begin to understand the question, a review of related topics in philosophy, mathematics, physics, biology, and other subjects are introduced through readings and a variety of classroom activities. A particular emphasis on teaching the mathematical roots of AI is made and includes the study of Hilbert, Russell, Church, Turing and others. Because we know that students learn best when the knowledge and experiences that each brings to the classroom are engaged[2], we begin the course with the activity shown in Figure 1. In this activity students attempt to answer a variety of questions related to philosophy of the mind and AI, and then share and debate them with their classmates. Many students are surprised that there are not simple answers to each question that can be memorized, but that instead each question requires further study to understand the issues related to each.

Two assignments in this unit are of particular importance. In the first assignment students explore web-pages and a variety of readings on Alan Turing in preparation for writing a paper on his role in the history of artificial intelligence and the debate over machine consciousness. In the second assignment students integrate everything that they have learned in the course with additional research to prepare for a formal debate on the question—"Is it possible for a machine to be conscious?" After a number of years of trial and error, we have found the following format to be most effective for stimulating purposeful debate.

Affirmative constructive speech (6-8 minutes)
Negative constructive speech (6-8 minutes)

Proceedings of the 2004 American Society for Engineering Education Annual Conference & Exposition

Break to coordinate rebuttal speeches (2 minutes)

Affirmative rebuttal speech (4-6 minutes)
Negative rebuttal speech (4-6 minutes)

Each side takes turns questioning the other side (30 minutes)

Try to answer the following questions and be prepared to discuss/debate them in class. Don't look up definitions in the dictionary or get frustrated if you find the task to be impossible at times—this is normal.

1. Observe your own thinking for several minutes (not easy). What was your mind doing during that time? What is thinking?

2. Observe your own consciousness for several minutes. What is it? Do you think animals are conscious? If yes, are all animals conscious, or is there a cutoff at some level of life forms?

3. Is it *theoretically* possible for a machine to think or be conscious? (Ignore limitations of technology, manufacturing, etc.)

4. Suppose a machine fools you into thinking that it is a human being. Does this mean that the machine is thinking or that it is conscious?

5. Can the workings of the human mind be written into a computer code (albeit a long one)?

6. Can a machine be creative? Have feelings? Have emotions? Have free will?

7. What is it that differentiates life from non-life? Can creatures that exist in a computer program be alive?

Figure 1 Assignment for bringing out student preconceptions about machine consciousness and other topics.

Turing Test

In Alan Turing's classic article "Computing Machinery and Intelligence",[3] he begins with the quote "I propose to consider the question, 'Can machines think?'" and then proceeds to describe a procedure for investigating machine intelligence. Turing writes:

> The new form of the problem can be described in terms of a game which we call the 'imitation game'. It is played with three people, a man (A), a woman (B), and an interrogator (C) who may be of either sex. The interrogator stays in a room apart from the other two. The object of the game for the interrogator is to determine which of the other two is the man and which is the woman. He knows them by labels X and Y, and at the end of the game he says either 'X is A and Y is B' or 'X is B and Y is A'. The interrogator is allowed to put questions to A and B thus:

> C: Will X please tell me the length of his or her hair?
>
> Now suppose X is actually A, then A must answer. It is A's object in the game to try and cause C to make the wrong identification. His answer might therefore be
>
> 'My hair is shingled, and the longest strands, are about nine inches long.'
>
> In order that tones of voice may not help the interrogator the answers should be written, or better still, typewritten. The ideal arrangement is to have a teleprinter communicating between the two rooms. Alternatively the question and answers can be repeated by an intermediary. The object of the game for the third player (B) is to help the interrogator. The best strategy for her is probably to give truthful answers. She can add such things as 'I am the woman, don't listen to him!' to her answers, but it will avail nothing as the man can make similar remarks.
>
> We now ask the question, 'What will happen when a machine takes the part of A in this game?' Will the interrogator decide wrongly as often when the game is played like this as he does when the game is played between a man and a woman? These questions replace our original, 'Can machines think?'

In this curriculum unit students learn about the Turing Test and participate in a number of activities that utilize the ideas behind it. These activities allow students to explore chatterbots—programs that attempt to simulate the conversation of a human being—while also developing critical thinking skills.

Students begin the unit by reading articles[4] describing and discussing the Turing test. Based upon these readings, students then work on an activity designed to bring out their preconceptions on the subject. In this activity students write down their best question for an interrogator to use to differentiate a human from a machine. Each of these questions is written on the board and the students then work together to decide which they think may be effective. For example, they may discuss if real world knowledge or semantics is necessary to provide a reasonable answer? Sample student questions are the following:

> If you were with 100% certain that you would die within a month, what would you do?
>
> Could you explain the meaning of ”Why did the chicken cross the road?”
>
> What question would you ask to test a computer’s intelligence? How would you answer it?

We have found the gender imitation game described by Turing to be an excellent classroom activity. In our game students in the class act as the interrogator and write questions that are delivered to a male and female located outside of the classroom. We have found that this activity is fun and very engaging for students, raises issues for discussion regarding gender and stereotypes, and requires students to practice critical thinking and group skills. During the activity much of the learning that takes place occurs as students work together to formulate questions. An example of the questions and answers in one class are shown in Figure 2. Once they have completed the game and attempt to identify the gender of the players, students continue their learning by reading the questions and answers recorded from previous games. From these games they identify good and bad questions and discuss the merits and flaws of each. They also conjecture on the gender make-up of the questioning team.

A student originally suggested this activity after having read Turing's paper and it has proven to be as engaging and complex as the issues it explores. Students are surprised and excited to explore gender issues in a computer science/technology class. We have several observations from watching the game. First, single sex questioning teams are at an extreme disadvantage and often cannot expose the imposter. Second, almost all teams ask questions based upon stereotypes, but then realize later that the answer tells them little. For example, one group asked the question, "How do you handle a flat tire?" Their discussion indicated that they expected the female to call for help and the male to put on the spare tire. As expected they received each of these answers ("Call AAA" and "Open the trunk, get out the jack and jack up the car, change tire!), but after much debate decided that the question was useless. Both a female that met their stereotype and a male imitating that stereotype would answer call AAA. Finally, just as in a Turing Test for differentiating humans from machines, the best questions require real world experience and more lengthy answers. For example, the question "What do you do in the morning?" resulted in responses of "Shower, wash face, brush teeth, put contacts in, get dressed, brush hair, put on mascara lip gloss, pack bag" and "I get up, get dressed and head over to breakfast." The second respondent was immediately exposed to be the imposter.

At this point the students have developed some sophistication in the relevant questioning and analytical skills and are ready to "converse" with chatterbots. In the first activity students converse with six different chatterbots. They are encouraged to ask questions inside and outside of the chatterbot's knowledge domain (if it has one) and probe its ability to parse sentences with varying syntax and ambiguities requiring semantics for understanding. They individually rank each program's humanness and list its strengths and weaknesses. Later this information is shared and discussed by the class.

Once students have completed the previous activity, they are then ready to begin a more in-depth investigation of chatterbots. In this assignment students choose one of the chatterbot programs that they have already worked with and write a paper describing in detail the logic it uses to imitate conversation. Key to this assignment is the student's effort to develop a logical plan of questioning that exposes the chatterbot's programming logic. For example, part of a line of questioning could be the following questions.

1. Are you alive?
2. Are you alive?
3. You are alive?
4. Are you alive.
5. Are you allive?

By comparing the answers to the first and second questions, the student gains clues into the randomness of the answers. Comparing the first and third questions provides information as to how the program deals with word order. The fourth question investigates the importance of punctuation and the fifth question explores the chatterbot's ability to deal with misspellings that a human could easily understand. Students are highly engaged in this activity and normally develop a sophisticated understanding of the workings of their chatterbot. One student mapped the chatterbot Alice[5] and posted his assignment paper on his personal website. He later received an e-mail from one of Alice's programmers (who had found his paper during a web-search) expressing amazement that a high school student had managed to so fully map Alice's logical structure.

What is your favorite scene in "Legends"?
X When the guy walks in and finds Brad Pitt and his brother's wife are hitting it off. It's an awkwa
moment for everyone.
Y Where Brad Pitt kissed his brother's fiance.

Name 7 fashion magazines
X Vogue, Bazaar, Elle, Mademoiselle
Y Seventeen, Cosmo, YM, Teen, Redbook

Where do you start shaving your legs?
X ankles
Y At the ankle

Describe the perfect date.
X Candle lit dinner followed by a walk on a secluded beach.
Y He has to be a guy that is nice enough to hold doors for me but not to egotistical so that he has t
for me.

What is a crimper?
X A crimper makes your hair wrinkly
Y A crimper crimps hair to make it look like how it does when you take you hair out of braids.

Did you like Topgun? What is your favorite scene?
X I love Top Gun. My favorite scene is where Tom Cruise is singing "You've Lost That Loving F
with his buddy in the bar.
Y I liked Top Gun but I last saw it a long time ago. The scene I liked best was where Tom Cruise
back with the lead actress at the end.

If you have red hair, would you wear a black or pink dress
X black
Y I'd chose the black dress if its not too skimpy.

How do you flirt?
X By smiling and acting very interested in what the other person has to say. Winking help too.
Y I'll look at a guy a certain way and talk to him when I see him.

Note: X is female and Y is male.

Figure 2 Example of Turing's imitation game. (Class was composed of 6 males and 4 females. In this class 7 students guessed the gender of X and Y correctly.)

Artificial Neural Networks
Artificial neural networks are designed, from a blue print of the brain, to simulate the brain's capability to think and learn through perception, reasoning, and interpretation. These networks use highly interconnected groups of neurons to process information in parallel. Although ANNs can learn through a variety of means, the most popular method is by example and repetition referred to as back propagation. In these networks the ANN learns to relate input and output variables through exposure to a training set consisting of input/output values. Once a network has trained on a data set, its knowledge can then be assessed and validated by testing its ability to predict output variables for a data set that it has never seen before. ANNs have a wide range of applications and details of their operation are described in many resources, such as J. Lawrence.[6]

Students are introduced to back propagation ANNs through directed readings and a variety of class activities. During this time they must complete three assignments in which they learn to use a feed-forward, back-propagation artificial neural network using *Brainmaker* software[7] to complete increasingly advanced, open-ended assignments. The first two of these assignments are based on tutorials that are provided with the software.

In the first assignment students become familiar with the software and begin to develop a feel for the behavior of ANNs. In this assignment they use an ANN to recognize fruit based upon information describing their shapes, sizes, and colors. An interesting aspect of this network is its use of color mixing. Fundamental to all phases of the assignment is to encourage students to experiment with the network to see how varying learning parameters, network architecture and training data affect network performance. In the first phase of the assignment, students learn how to train, test and validate an ANN by using data that is provided to them. In the second phase of the assignment, students learn how to work with data files by adding additional facts to the training set. For example, the original data set includes only examples of bananas that are yellow, but they may be green also. In the third phase of the assignments students modify the dataset and the input/output structure of the network architecture by adding additional fruits and characteristics (such as texture) to the training set. In the final phase of the assignment students make up their own network from scratch to model a different subject.

In the second assignment students use an ANN to find a relationship between housing parameters (living area, number of bathrooms, etc.) and selling price. Again a data set for training the network is provided. The objective of this tutorial is to expose students to more a more realistic training situation and to techniques for understanding the trained network through a sensitivity analysis.

In the third assignment students in the class must work together to collect a dataset to be modeled. For example, in one year the class polled the student body for their political preferences in an upcoming presidential election. They used an ANN to model the relationship between the information about the survey respondent (sex, age, and other variables) and their preferred candidate. From the trained model students could then predict candidate preferences and examine which variables are related to them. In addition to applying what they have learned in the two previous assignments, students learn about sampling and data collection issues. They also must decide together upon a topic that is worthy of study and appropriate for ANN modeling. During this process students discuss and debate the quantity and quality of the data that they will collect. Will it be biased or independent? Will there be enough data to train a network? During the process of choosing a topic, collecting data, fitting a model, and analyzing the results, students encounter and learn about many of the potential pitfalls that they will face in their independent research project.

Independent Research Projects

The second half of the course is focused on completing an independent research project that involves the use of an ANN. This phase of the course is structured to simulate a professional research experience in which the instructor plays the role of a funding program director. The steps are as follows:

1. Students receive a request for proposals (RFP). To add realism, the RFP is written in a style similar to research funding agencies.
2. Students devise an original research idea and discuss it with the program director.
3. Students submit a formal research proposal that communicates what they want to accomplish, justifies the value of the research, and devises a detailed plan of action showing that the research can be successfully completed.
4. The proposal is peer-reviewed by classmates and other qualified reviewers if applicable. The final decision for acceptance rests with the program director, but it is based upon the peer reviews.
5. Several times during the research period students informally present their research to their classmates and receive feedback and advice. They are also required to submit a short progress report midway through the semester.
6. At the end of the semester students submit a research paper that is peer-reviewed by their classmates and make a formal presentation.

Throughout this process the importance of helping their classmates through peer-review and more informal means is emphasized and encouraged. The following are summaries of three students projects.

Modeling student admissions

In this project an 11th grade student used an ANN to model the admissions decision process at St. Paul's School. Similar to the college admission process, committees read through application files and decide whether to accept, deny or waitlist an applicant. The student was granted access to student admissions files (with identifying information deleted). Using these files, he developed an ANN that used the geographic location of the applicant, grade applied to, sex, minority classification, Secondary School Admission Test (SSAT) scores (verbal, quantitative, and reading comprehension), interview score (academic and personal) and the number of the applicant's alumni ancestors as input variables. The out variables were whether to admit, waitlist or reject the applicant. The student used some of the available data to train the neural network and then set aside the rest to validate it. Using the validation set, he found that the network made the same admissions decision as the admissions office 77% of the time.

Of particular interest in this project was a sensitivity analysis that showed which variables were most important in the admissions process. This was not only interesting from an academic standpoint, but also useful to the school which was under pressure from alumni who expressed concerns that alumni relationships were not being given proper consideration in the admissions process. The ANN showed that having alumni relations was the most important factor in the admissions process. Other factors that were found to be important were interview ratings, standardized test scores and minority classification.

Modeling Water Disinfection
In this project a student used an ANN to model data from an experimental study[8] on the use of ozone to inactivate a parvovirus in a synthetic and an actual industrial water source. The goal of his analysis was to predict the necessary ozone dose needed to disinfect the water as a function of specific environmental conditions. The network consisted of six inputs (time, alkalinity, organic carbon concentration, initial virus concentration, sonication, and ozone residual) and one output (virus concentration). The network was found to be highly effective in predicting the outcome of ozone disinfection. A sensitivity analysis revealed that the network learned relationships among the variables similar to known and accepted trends.

It was shown in this project that the ANN modeled the ozone disinfection process more accurately than accepted Environmental Protection Agency (EPA) and conventional modeling techniques. This pointed to the potential for using ANNs to increase the efficiency of the disinfection process. Also, from the sensitivity analyses, optimum ozonation levels and environmental situations were quantified. The results of this project were presented at the *Florida AI International Conference*[9] and published in *Environmental Engineering Science*[10].

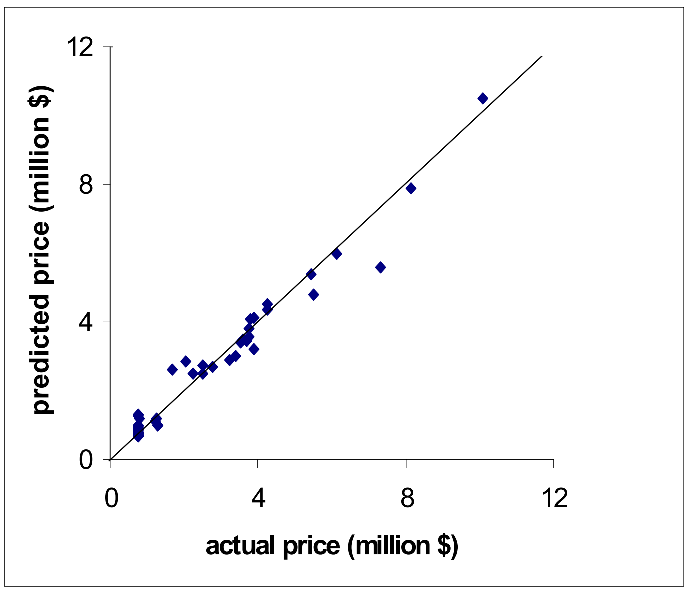

Figure 3 Comparison of ANN prediction and actual list prices of Greenwich houses.

Modeling Real Estate Prices
In this project a student used an ANN to model housing prices in his hometown of Greenwich, Connecticut. Input variables to the network included house size (square ft., number of bathrooms, number of bedrooms, number of garages), style (3 styles), land (acreage, pool, tennis courts, lake frontage, ocean frontage) and location (9 neighborhoods). The output variable was

the list price of the house. The excellent ability of the network to predict housing prices is illustrated in Figure 3. Through a sensitivity analysis the student investigated the effect of various input variables on the listing price.

In his project report the student noted the potential uses of his network. They included assessing the effect of input variables on price (i.e. predicting the value that adding a bathroom will have on the house price or measuring the value of water frontage), detecting price trends in the market, and helping determine appraisal values.

ASSESSMENT

Grading in the course was largely based on student performance on papers and presentations. During the research phase of the course, the effort and quality of peer reviews and other means for assisting classmates was included in addition to individual project performance. A mid-year exam was administered at the end of the fall semester (see Fig. 4). Because students were typically highly engaged and motivated in the class, they tended to score well with an average score of about 90%.

Student interest in the course was tremendous. When it was first offered at St. Paul's School, only three students elected to take the course. Two years later 38 students (about 30% of the graduating class) signed up for the course even though the maximum class size was only 12. 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 instead of tests, the flexibility to explore topics that interested them in depth, and the opportunity to conduct independent research. The school administration also viewed the course positively and highlighted it in the recruitment literature.

One other measure for assessing the course is the number of students formally continuing their research after the course ended. While seniors did not continue their research projects during the following year due to graduation and making the transition to college, most underclass students in the course continued their research the following year. Their accomplishments included a refereed article in a professional journal, a presentation at a professional conference, and several regional and national awards in science fairs.

DISCUSSION

Although this paper describes how the curriculum can be implemented as a full-year course, the units and activities are modular and can thus be used in a variety of classroom settings. For example, a course in robotics could include the machine consciousness unit. In Andam et al.[11] we describe how elements of the Turing test curriculum were used in a workshop to help prepare pre-service middle and high school teachers to integrate technology and engineering into their future classrooms. The workshop was required of all Smith College education graduate students and thus was populated by students with varying content specializations. In the workshop the students first worked through the Turing test activities themselves, and then they used the activities to teach a class of middle school girls. Although the field of artificial intelligence and the idea of teaching a class in the subject initially intimidated the education students, their student teaching experience was very successful and highlighted the potential use of AI as an interdisciplinary instructional tool. It is interesting to note that each of the teacher's different content knowledge enriched the classroom teaching. For example, the pre-service English teacher led a discussion about the importance of semantics in understanding language and the mathematics teacher used the activities to present topics in logic.

We feel that part of the success of the AI course results from our intended learning outcomes differing from those in typical high school courses. The development of critical thinking skills is an important learning outcome in our course and many of the activities are intended to support their development. Helping students make connections among the subjects that they learn in school is also an important learning outcome. As discussed in Ellis et al.,[12] we typically leave it to students to make these connections and see the big picture. The result is that they are often unsuccessful and their knowledge is domain specific. The interdisciplinary nature of AI helps make these connections explicit through application. For example, during a discussion of artificial life the students became confused. This became an important teaching moment when it was revealed that the confusion was based upon the students' misconception that the laws of physics did not apply to biological systems. A third intended learning outcome is the development of metacognitive skills. For example, in the research section of the course students learn to set their own research goals, plan of action and timeline. They then monitor their progress toward meeting their goals. The peer review process also provides practice for developing assessment skills. Finally, because ANNs are founded in the study of the brain, parallels to human learning exist and are explicit in the field's terminology. The study of ANNs offers students the opportunity to see, manipulate, and measure a form of learning that takes place.

CONCLUSION

We have developed curriculum units for teaching AI that have been shown to engage high school students. A key to their success has been to match the curriculum content and intended learning outcomes with the needs of students. By making an exploration of the possibilities (both theoretical and technological) of an artificial brain the unifying principle of the course, the content is engaging to students at a time when they are developing their own identity. Because the intended learning outcomes focus on thinking skills and knowledge integration, they address student needs that are often not met in other courses.

ACKNOWLEDGEMENTS

This research has been partially supported with generous grants from the GE Foundation, the Ford Motor Company and the National Science Foundation (grant BEE-0230625).

Artificial Intelligence Mid-Year Exam

1a. What is Godel's incompleteness theorem and how is it used by Penrose in the debate over machine consciousness?
1b. What are the four major positions on machine consciousness?

2a. How do classical AI and connectionism differ?
2b. In what ways is connectionism superior to classical AI?
2c. In what ways is classical AI superior to connectionism?

3. Write down a strategy for exposing a computer in a Turing test. Include multiple attack strategies and example questions for each one.

4a. What is a Turing machine? What is its significance in the AI movement?
4b. What is a universal Turing machine? What is its significance in the AI movement?
4c. Write a Turing machine code that given any tape will switch the 0's to 1's and the 1's to 0's. It will continue until it reads two consecutive blanks when moving to the right.

5. Briefly discuss how the history of classical AI is both a success and a failure. Give specific examples to support your opinions.

6a. What is overtraining in an artificial neural network? Give at least two reasons why overtraining might occur? Explain for each one how it contributes to overtraining.
6b. Explain why each of the following might occur for a neural network.
- low training error with low validation error
- low training error with high validation error
- high training and validation error

7. An ANN is being trained to differentiate between the letters P and Q. A database of handwriting samples of the letters is being collected to train the network. There are 50 inputs, 1 output (0 for P, 1 for Q), 8000 training facts, and 1000 validation facts. The result of training with differing number of hidden nodes is shown below.

Case number	Number of hidden nodes	# of iterations	Training error	Testing error
1	10	100	21.2	11.1
2	10	500	9.3	8.3
3	10	1000	7.2	9.0
4	10	5000	6.1	9.9
5	10	10000	5.1	17.5

a. What errors could the investigators make when collecting the data set? How would each of these errors affect the ability of the network to correctly learn character recognition?
b. Will the training error eventually go to zero? Explain very briefly.
c. Between what iteration numbers is the optimal network reached?

8. Three synapses lead into a neuron that utilizes a sigmoid function. The weights are 2.0, 1.0, and –1.0 respectively. The inputs from other neurons are 2.0, -4.0, and 2.0 respectively. The bias to the neuron is –5.0.
a. What is the output of the neuron?
b. If the neuron is not fully trained, which of the above numbers will change during training?

Figure 4 Example mid-year exam.

Proceedings of the 2004 American Society for Engineering Education Annual Conference & Exposition

REFERENCES

[1] M. Minsky, *The Society of Mind,* Simon and Schuster, New York (1986).

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

[3] A.M. Turing, "Computing, Machinery and Intelligence," *Mind*, Volume 59 (1950).

[4] R.A. Kurzweil, The Age of Intelligent Machines, MIT Press, Cambridge, MA (1992).

[5] *A.L.I.C.E AI Foundation,* <http://www.alicebot.org/> January 13 (2004).

[6] J. Lawrence, *Introduction to Neural Networks*, 5th ed. California Scientific Software, Nevada City, CA (1993).

[7] California Scientific Software, *Brain-MakerProfessional: User's Guide and Reference Manual*, 4th ed., Nevada City, CA (1998).

[8] K. Flood, *Inactivation of parvovirus*. MS thesis submitted to the University of Connecticut (1995).

[9] S. Heck, G.W. Ellis and D. Grasso, "Artificial Neural Network Analysis of Ozone as a Water Disinfectant," *Florida AI International Conference*, May 11-15 (1997).

[10] S. Heck, , G.W. Ellis, and V. Hoermann, "Modeling the Effectiveness of Ozone as a Water Disinfectant using an Artificial Neural Network," *Environmental Engineering Science,* Vol. 18, No.3 (2001).

[11] G.W. Ellis and B. Andam, "*Teaching Teachers to Teach Engineering", 2004 American Society for Engineering Education Annual Conference & Exposition,* Salt Lake City, Utah (submitted for purblication

[12] G.W. Ellis, B. Mikic and A.N. Rudnitsky, "Getting the 'big picture' in engineering: Using narratives and conceptual maps," Proceeding of 2003 American Society for Engineeirng Education Annual Conference and Exposition, Nashville, Tennessee, June 22-25 (2003).

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 has 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.

BAABA ANDAM is a sophomore engineering science student in the Picker Engineering Program at Smith College. She has been working in engineering education research for one year.