

MODERN DSP EDUCATION: THE STORY OF THREE GREEK PHILOSOPHERS¹

Joseph Picone, Jonathan E. Hamaker, Robert Brown²

Institute for Signal and Information Processing
Mississippi State University
Mississippi State, Mississippi 39762
email: {picone, hamaker, brown}@isip.msstate.edu

John H. L. Hansen and Ronald A. Cole

Center for Spoken Language Understanding
University of Colorado, Boulder
Boulder, Colorado, 80309-0258
email: {jhlh, cole}@cslu.colorado.edu

ABSTRACT³

In many ways, education has come full circle in little more than 2,500 years. From the great Greek philosophers, we learned the importance of asking questions, searching for a deep understanding of the most abstract of concepts, and expressing such concepts in rigorous mathematical form. Since then, we have witnessed an endless struggle over what to teach, how to teach it, and when to teach it. In the meantime, the body of foundational scientific knowledge we seek to understand continues to grow at an astonishing pace. Today we find that most students taking DSP do not possess traditional signal processing fundamentals. In fact, teaching DSP is one big circular argument. Only when you reach the end of the course are you able to understand how to convert an analog signal to a digital signal (because inside a state-of-the-art converter is an embedded DSP system).

Fortunately, the Internet is enabling a revolution in the educational process. Students now have access to a wealth of information, and can control its flow. How, then, do we harness this vast expanse to accelerate the pace of learning? How also should this access to knowledge be structured in a classroom setting to be most effective in the learning process? Though it is premature to provide definitive answers to this question, we can describe some evolving directions in Internet-based education that provide insights about where things are going in the next five years. Certainly, Java and the Internet II are central to this strategy. Assuming intellectual property concerns can be resolved, the notion of printed textbooks and accompanying software are giving way to electronic “live” documents and open source toolkits. We present a vision of how such technology can enable students to learn in non-traditional manners, and discuss why it is important that we accommodate this learning style.

-
1. This work was supported in part by the National Science Foundation under Grant No. EIA9809300.9811040.
 2. Currently with Digidesign, Inc., 3401-A Hillview Avenue, Palo Alto, CA 94304-1348.
 3. This paper is better read on-line so you can easily explore the links. An on-line copy of this paper is available at:
http://www.isip.msstate.edu/publications/journals/ieee_spmag/1999/education/.

1. INTRODUCTION

The classroom dynamic has changed — if we are certain of one thing in this changing world, it is this. Ask a student to do a literature search. The enterprising student will return with the results of a web search. If it is an obscure topic, the student will perhaps complain that he or she couldn't find anything really useful on the web, what should they do? Ask your students how many of them have read the textbook they bought for a course. Many will reply that they don't read the textbook as carefully as they should because it is too difficult to follow (and they seem to say this about any textbook you choose). Occasionally you might hear a truly honest response: “Why would I open it? That would only lessen its resale value.”

Are students less capable today in their analytic and mathematical skills than in years gone by? There does seem to be a wealth of evidence that students today are poorly prepared for college [1] and the source of the problem is K-12 education [2]. Yet, despite this perceived drop in students' skills, technology is certainly progressing at a dizzying rate. In this paper, we adopt a point of view that students are not less capable, it is simply that there have been dramatic changes in what, when, and how students need to learn. We are no longer afforded the luxury of teaching in a linear fashion: algebra, calculus, differential equations, discrete mathematics, circuits, signals and systems, complex variables, and then, finally, digital signal processing (DSP). The solution is a more random approach to learning sometimes referred to as “learning on-demand” with the Internet as the great enabler for this revolutionary change in education.

Of course, this is nothing new. A long time ago, Greek philosophers grappled with this very question. From **Diogenes** (413-323 B.C.), we learned the importance of learning by doing: “virtue was better revealed in action.” From **Socrates** (469-399 B.C.), we learned the principle of life-long learning: “I know of nothing except my own ignorance.” From **Plato** (429-347 B.C.), we learned the importance of mathematical modeling: “the reality which scientific thought is seeking must be expressible in mathematical terms.” Perhaps this was the beginning of DSP?

Advances in interactive language technologies will eventually revolutionize learning and training. Computer-generated animated characters with expert domain knowledge will interact with learners much like the best teachers do today. For example, below is one such scenario that is in the preliminary research stage [3]:

Julia's ninth grade conceptual physics class includes a learning module on the nature of time. Julia consults her favorite animated agent, Teacher Molly, and explains her assignment. Saint Augustine appears, introduces himself, and asks: “What is time?” He proceeds to describe the puzzle of time. After an interesting discussion, Albert Einstein arrives, and provides some excellent description of the relativity of time, with interesting animation showing the relationship between travel and aging. During the discussion with Dr. Einstein, Dr. Stephen Hawking interrupts to point out the three arrows of time. Excited by these mysteries, Julia asks Teacher Molly how one of the arrows, entropy, can be true in an organized world. Teacher Molly suggests that Julia use the “active worlds” program to simulate a hypothetical world. After several simulations, and some discussions with Teacher Molly (who has the annoying habit of making Julia answer her own questions), Julia understands the trade-off between entropy and biological evolution. For her science project, Julia incorporates her simulation into a discussion between Albert Einstein and Charles Darwin, which is made available to other students.

We can only imagine how learning will change when students have access to all available knowledge (and frequent access to the experts who created it). Learning will be interactive, individualized, self-paced and infinitely variable. The big question is: How far away is “eventually?”

What are the major barriers preventing colleges and universities from developing curricula using such powerful tools? Beyond the fact that adequate core technology still does not exist for many of the previously described functionalities, such authoring environments require significant multidisciplinary expertise and collaboration. There are numerous legal issues involved in the copyrights and licenses potentially required for such technology (though we will focus here on open source initiatives). Perhaps most importantly, the traditional providers of content, textbook manufacturers, do not see a clear path to a revenue stream with electronic publishing. Hence, the intellectual and investment capital that feeds the publishing industry is only just beginning to fuel Internet-based development of educational materials.

In this paper, we review what progress has been made recently with Internet-based educational technology and tools. Now that the enthusiasm for some new development environments has subsided, the real work is beginning as we realize technology without pedagogy amounts to a lot of wasted programming. Next, we discuss the drawbacks of some of these approaches, and focus on the next steps being taken to deliver a comprehensive range of tools to support learning at all levels of a discipline. Finally, we turn towards the future and speculate on what forms educational materials might take if human language technology delivers even half of what is promised.

2. IN THE BEGINNING, THERE WAS JAVA

The goals for any engineering course are to convey to student engineers a broad understanding of the mathematical theories underlying engineering concepts and to develop engineers who are able to apply those concepts with ingenuity and expediency. The second of these skills must be learned by repeated application of the concepts over a wide variety of test conditions. In doing so, the student builds an intuition for the problems involved. When teaching science or engineering, it is clear that students learn in different ways. Some prefer to focus on mathematical structure and understand the balance between “input” and “output” characteristics using mathematics. Others comprehend concepts more in a visual manner, where illustrations visualization can reinforce complex concepts. Still others prefer to “get in the lab” and try out either a workbench circuit simulation or a computer simulation of the concept. No one method is right or wrong, and teachers need to recognize that more than one route may be needed to reach their students.

For many years, the primary method of teaching engineering concepts has come in the form of classroom interaction and homework exercises. These methods do an adequate job of conveying mathematics as well as the mechanics of solving problems with a predefined structure or a “recipe,” but often leave the student lacking the intuition necessary to apply the concepts to new problems without the same structure. As an example, consider the amount of time during an open book exam students spend hunting in their notes for an equivalent problem to pattern or “plug-and-crank” a set of equations for a solution? If students truly understood the concepts, they would be able to extend their knowledge to address new problems they have not seen before.

The main stumbling block in building intuition about concepts is the student's inability to visualize the items being learned or their unwillingness to commit the necessary time to understanding the problem or concept. Many concepts have complex visual representations that cannot be sufficiently explored through the standard classroom experience. One example, shown in Figure 1, is the relationship between the poles and zeros of a system and the three-dimensional surface representing the system's response. This phenomenon is usually sparsely covered in textbook problems if not ignored altogether. Yet, it is an important concept in the understanding of linear systems.

The introduction of computer-based visualization into the educational process has been seriously pursued for some time [4]. Perhaps the most striking example of this is the MathWorks Inc. simulation tool known as Matlab [5] — by far the dominant software tool in DSP today. In fact, many undergraduate textbooks now integrate Matlab exercises and examples into their text. Student editions of Matlab are inexpensive and bundled with entry-level textbooks. Nevertheless, such tools have two major drawbacks: a steep learning curve and the necessity of having access to a commercial product in order to test a specific application. There are a number of public domain Matlab-like alternatives [6, 7], and some alternative scripting environments [8], but none that have enjoyed as much fanfare as Java [9].

A good example of a DSP concept that can be greatly enhanced through a Java applet is graphical convolution — an algorithm used to compute the time-domain response of a linear system to a signal. An applet implementing graphical convolution [10, 11] is shown in Figure 2. With the aid of computers, convolution is a fairly easy concept to teach. Prior to computers in the classroom, convolution was an extremely tedious topic to teach mainly because of the lack of a good animation facility. With this Java applet, students can investigate pre-stored waveforms, or draw their own waveforms. By pressing the convolve button in the upper left of the screen, students can see the result unfold via a Java-based animation. The convolution tool allows students to easily solve textbook problems, and lets them develop intuition about convolution by trying lots of examples and rapidly visualizing solutions to incremental variations of a classroom example. The bottom line is that students can **interact** with the algorithm.

But what is the principal advantage of a Java version of this tool? Certainly, such animated tools have existed long before the advent of Java — many such signal processing tools were very nicely implemented on Apple's Macintosh in its heyday. The answer lies in issues related to ease of use and portability. Students need not invest much time downloading and installing the tool. They simply click and run. For students that have minimal investment in signal processing, such as entry-level undergraduates or electrical or computer engineers not specializing in signal processing⁴, a point and click web-based tool is far better than any tool requiring installation of an environment, local infrastructure, or moderate amounts of programming. Java is attractive because the code is free and is inherently platform-independent⁵. The Java approach allows a tool

4. This is primarily true for academic programs that do not force students to do a lot of Matlab prior to upper-level engineering courses. Some programs that now introduce Matlab at the entry-level have the luxury of dealing with students well-trained in the art of computer simulation in the upper-level engineering courses.

5. In theory Java is platform-independent. In practice, the road has been bumpy, and support of new versions of Java by the major browser providers has been less than ideal lately.

to be accessed by the widest audience possible.

Another attempt to develop such an interactive tool is shown in Figure 1. In this applet [12], users can create poles and zeros, move them around in the s-plane or z-plane with the mouse, and see the resulting changes in the frequency response of the associated system transfer function. This is done interactively — the plots are updated continuously as the user changes the parameters of the system. You can also generate a 3D view of the magnitude response of the system, and inspect this plot by rotating it using the mouse. This applet attempts to teach students the relationship between frequency response and poles and zeros. An extensive on-line tutorial is included that walks users through the use of the tool and elementary concepts.

It would be nice if the story ended here. The “one applet per concept” paradigm seemed like an elegant solution if only we had an authoring environment that would make development of such things fairly easy. Unfortunately, as we will see in the next section, presentation of a tool doesn't necessarily mean learning will occur. Moreover, growing pains with the Java environment have resulted in numerous technical difficulties in the classroom (the most common problem involves a lack of backward compatibility of newer versions of Java with older versions of the browsers).

3. “A FUNNY THING HAPPENED ON THE WAY TO THE FORUM”

The convolution demonstration was introduced into the classroom for the first time in 1996. Since then, it has been used regularly in an entry-level Signals and Systems courses at Mississippi State University. It is interesting to note that by the end of the semester such tools are viewed very positively by students, and preferred over traditional textbooks, homework solutions and such. We might ask ourselves ‘Is this just a novelty?’ or are students really learning more or better from such experiences? What has been most disappointing in practice, however, is that most students will not even open the applet during the semester unless they are “motivated” to do so. There are many reasons for this, and we don't want to digress into a lengthy discourse on education and societal trends, but we can simply state our observation that motivating the student to use such tools is as critical as supplying technically sound tools. If, in the end, all we have achieved is another method to make students respond to artificial stimuli such as an exam deadline, then we haven't achieved our true goal of transforming them into life-long learners. The lesson here was simple: our “build it and they will come” approach failed.

There are some interesting trends developing that involve the integration of software and data into classroom learning. In many ways, these trends are synergistic with a rapidly growing open source [13] movement focused on placing technology on the web in ways that make it easy for people to collaborate and share ideas. Some research areas within the digital signal processing community, for example human language engineering [14], are promoting this strategy for the development of new multidisciplinary educational programs. One of the first comprehensive surveys of human language technology [15], funded by a number of interests including NSF and the European Union, was published here as an electronic document. The delivery of this document is in a fairly conventional hypertext form. However, it is packaged with other web resources that enable one to explore speech technology.

One of the first journals to consider integrated data such as speech and audio samples was

Elsevier's Speech Communication, thanks in large part to the tireless efforts of Christian Benoit. Other journals in the signal processing area have been slow to respond, but great progress has been made with the IEEE Signal Processing Society web site at Rice University [16]. It seems natural that for students learning about techniques in speech coding, synthesis, or enhancement, that they should have access to the processed sounds. It is somewhat difficult to truly explain what perceptual processing artifacts result when we reduce the bit representation for a set of model parameters, or introduce a new way to characterize the speech spectrum for background noise suppression. The only way for students to understand many signal processing algorithms which are used for speech or audio applications, is for them to be able to listen to the processed signals. In many respects, the cornerstone paper by Klatt [17] was instrumental in introducing a whole generation of students to the field of speech synthesis, not only for its historical treatment of the subject, but more so because it was one of the first journal publications which included a recording of example synthesized speech signals from more than 40 years of research in the field.

A somewhat related activity also centered at this site was the Free Speech journal [18] — an attempt to deliver a peer-reviewed research journal entirely electronically. This was subsequently subsumed by the Elsevier Science's publication Speech Communication [19]. This journal is attempting to provide a forum that supports dissemination of other forms of data related to a technical publication beyond the document itself (for example, processed speech data). There are compelling reasons for such a forum, including verification of conclusions drawn by an author, and direct comparisons of new approaches on the same data. Not surprisingly, other organizations, such as the IEEE, are contemplating such directions.

The above web sites, however, still represent a rather conventional delivery of material. Though users can traverse the information in a nonlinear, or random fashion, the material is still largely non-interactive. A better example of the wave of the future in signal processing education is the CSLU Toolkit [20]. In addition to the distribution of software and technology, this toolkit includes numerous on-line demonstrations of the underlying technology. For example, a spectrogram reading tutorial [21] is available that integrates audio-visual information to teach students the fundamentals of speech production. Another interesting tutorial [22] teaches users how to build a neural network-based speech recognition system using the toolkit. An example of this tutorial is shown in Figure 3. The key point here is that this toolkit is available for many common computing platforms, and is being actively developed as an educational tool for a wide range of activities ranging from K-12 education to graduate studies.

Another more focused project involves the development of public domain speech recognition technology [23]. This project, funded in part by the National Science Foundation, has three important components: software, education and support. In addition to the distribution of a speech recognition toolkit, this site features tutorial materials centered around the components of this technology, documents describing key algorithms in great detail, and a collection of pointers to other resources on the web. An example of the educational resources available at this site is shown in Figure 4. Other resources, such as on-line support, application-specific toolkits, and supporting software are also available at this site. A series of training and planning workshops will be held annually to promote education and dissemination of information related to the project.

A novel aspect of this site is a remote job submission facility [24] which is shown in Figure 5.

This site is being developed in anticipation of the next generation of the Internet where connection speeds will reach into the Gbps range, and new capabilities for resource sharing will be introduced. At this site, users can submit a speech recognition job over the web, using user-supplied data or pre-stored data. The system returns results at varying levels of detail. One use for this site is to compare a user's algorithm to a standard algorithm on the same data. Another function is to allow users to learn how the technology works without having to perform a complicated installation on their local machine. The source code for this system is distributed via an anonymous concurrent versions system (CVS) server [25] that, if needed, lets users easily maintain a parallel copy of the software that exists at this web site.

A final example of the power of the web is an interesting text-only interface involving the use of a relational database to provide information on terminology [26] as shown in Figure 6. In this system, it is possible to search for the meaning of a term in many different ways, and to easily learn about related terms. Each term in the database is encoded with a set of attributes from which all other information can be automatically derived. Such interfaces, though presently constructed manually, are the topic of a great deal of research in natural language and machine translation communities. Providing structure to data automatically is still a great challenge. Web search engines are a good example of how far we have come, and yet also an example of how far we have to go to accomplish useful semantic queries of information. Yet, from such interfaces it is clear that web-based documents offer a potential that vastly exceeds printed documents.

As interesting as this may sound, most of what we have discussed thus far does not address the motivational issues so crucial to modern education. Very little effort has been spent at the above web sites on how to present the material. Most of the associated documents have a standard engineering look and feel to them. Random access of this information is difficult — anyone who has tried to use web search engines can testify to how hard it is to sort through the vast amounts of information available on the Internet. In the next section, we look at some interesting paradigms for presentation of information that attempt to go far beyond the web as we see it today.

4. BUT HAL 9000 WAS JUST A VOICE...

A powerful metaphor for intelligent computing has been HAL 9000, the computer featured in the movie titled *2001: A Space Odyssey*. HAL displayed vast amounts of intelligence and humanlike emotions. The subtle and complex man-machine interactions depicted in this movie have remained cultural icons for three decades. However, a vision of man-machine interfaces is evolving that involves animated computers. These are computer interfaces that recognize, synthesize, and integrate audio and visual information, such as vocal emotion, body language, and facial expressions. These were formerly thought to be extraneous to the man-machine communication process, and now are recognized as extremely important in understanding the intention of the user. HAL was just a voice, and technology finally appears to be moving beyond simple audio and keyboard interfaces to communicative agents that simulate face-to-face conversational behaviors.

An early approach to developing interactive educational materials was the use of virtual reality engines to let the student explore a 3D world on-line. A good example of this is the National Geographic web site [27] in which one can explore the solar system. Similar sites exist that let

you explore products you wish to purchase, manipulate robots, etc. One of the more exciting areas of research in this direction involves the use of tactile feedback and other senses to provide users the ability to experience an object or environment remotely. While the technology has not progressed yet to the point where such things are possible with the realism desired by a generation weaned on video games, clearly such applications offer unique and interesting possibilities. A good example is a project spearheaded by several museums that will let users handle rare museum artifacts over the web using data gloves. This is an interesting application because it allows users to do something they could never do at the museum — touch a rare artifact. The potential for interactive educational agents will only be achieved when these applications offer users experiences that cannot be easily accessed in the real world.

Enhancement of the traditional computer interface with two-way audio and visual information is one area of research receiving great attention these days, particularly in the Signal Processing Society. Animated conversational agents, such as those described in [3, 29] are being developed that recognize and synthesize paralinguistic information in an attempt to humanize the man-machine interface. One example of this currently under research is Baldi [3], shown in Figure 7. Baldi combines audio and visual features during speech synthesis. These features are often complementary and add redundancy that humans require to maintain conversations without feeling unduly burdened. We effortlessly switch between these cues as needed. Interfaces that are ultimately as powerful as those shown in Figure 7, and can be configured to suit profiles of users, offer unlimited potential in the development of systems that provide information in many different learning styles. This seems to be crucial because it is clear that there is no best learning style, and that learning is most efficient when the presentation matches a learning style. Computer technology offers great promise of customization through the classroom — with a teacher becoming more of a roving monitor rather than the primary information provider.

We as signal processing engineers also have a responsibility to facilitate the evolution of technology to provide innovative ways for teachers in other fields as well. One project which is beginning this year is an effort by researchers at Michigan State University (ECE, Humanities), University of Colorado (RSPL-CSLU), and University of Pennsylvania (LDC) to transition the enormous audio corpus which is the Vincent Voice Library (60,000 hours of historical audio recordings) to a “National Gallery of the Spoken Word” [30]. Such an effort requires cataloging, organizing, and providing meta-data information, as well as the development of an effective phrase search engine for the WWW with data compression and watermarking technology. Such a system will allow students learning about history, politics, or any subject which might have been represented with an audio recording over the past 50 years, to get instant access to the people who “talked the talk!” (as opposed to hearing the teacher say, I guess you just had to be there). How many future astronauts might be inspired after listening to Neil Armstrong say “One small step for man, one giant leap for mankind.” So, while more effective ways of teaching DSP will come from better interactive and portable learning tools, it is important to recognize that this will also benefit other disciplines in a way we could not have imagine.

A final example of some evolving technology that is becoming increasingly relevant to education is the 3D immersive environment known as a CAVE [31]. An example of such a system currently deployed at Mississippi State’s Engineering Research Center is shown in Figure 8. In these environments, students can explore solutions to problems and walk through the problem space

with an amazing amount of realism. Groups of students can also work together in such environments. While the cost of such facilities is presently prohibitive, these systems give us a glimpse of what education might look like in the future. Integrating immersive visual and tactile environments with audio provides a powerful environment in which learning by example can take place. The impact of such environments on learning is hardly understood and certain to be the topic of great debate in the next millennium. However, it is clear that students accustomed to such technology will no longer be willing to sit still for chalkboard lectures on digital signal processing theory.

5. SUMMARY

We have explored the relationship between DSP education and the Internet, and shown some interesting trends involving platform-independent programming paradigms based on Java, open source software, and interactive immersive environments. We mainly addressed technologies we believe will be deployed within the next five years. We did not attempt to grapple with the shape and form of education 50 or 100 years from now, for we are certain there will be a holographic communicative agent named Socrates running around your living room imploring you to seek a deeper understanding of the most fundamental aspects of human existence.

Ironically, despite the introduction of such fascinating technology, it seems to be the human element that is so important to the learning process. DSP, like most sciences these days, can at the surface appear to be simply a pile of obscure mathematics, or at the deepest levels a bundle of details and heuristic algorithms with no elegant structure. To make matters worse, students are being introduced to the subject earlier than ever in their education. The instructor becomes essential in such a situation, because this person must guide these students through their journey by tutoring them on the art of learning (rather than focusing solely on transmission of factual information from one to the other). It is perhaps this aspect of scholarship that, to this day, is most impressive about the early Greek philosophers — they seemed to deeply understand the essence of learning and a process by which it could be achieved.

In any paper like this it is hard to do justice to all of the relevant research. Rather than attempt a comprehensive overview, we have instead preferred to provide a peek into a new world of interactive learning tools. To probe further, there are several projects of note. It seems like kids have all the fun these days, and [33,34] are representative of some typical directions in K-12 education. There are also a number sites devoted to providing searchable databases of learning technology resources [35-37]. Tools to help automate the process of learning over the Internet are arriving, and WebCT [38] is a good example of a tool gaining widespread acceptance at the university. Specific to DSP, the Matlab site [5] is a definitive resource for any engineer. An interesting new site that has just come on-line is dspGuru [39] that features a documentation of DSP tricks of the trade. Finally, Tech On-Line [40] offers many interesting capabilities, including the ability to write code, download it to a DSP board, and run it, all from the web.

REFERENCES

- [1] *A Nation At Risk: The Imperative For Educational Reform*, The Commission on Excellence in Education, Washington D.C., USA, 1983.
- [2] *Report to the President on the Use of Technology to Strengthen K-12 Education in the United States*, The President's Committee of Advisors on Science and Technology, March 1997.
- [3] R. Cole, D.W. Massaro, J. de Villiers, B. Rundle, K. Shobaki, J. Wouters, M. Cohen, J. Beskow, P. Stone, P. Connors, A. Tarachow, and D. Solcher, "New tools for interactive speech and language training: Using animated conversational agents in the classrooms of profoundly deaf children," presented at the ESCA/SOCRATES Workshop on Method and Tool Innovations for Speech Science Education, London, UK, April 1999.
- [4] R.J. Moorhead II and Z. Zhu, "Signal Processing Aspects of Scientific Visualization," *IEEE Signal Processing Magazine*, vol. 12, no. 5, pp. 20-41, September 1995.
- [5] <http://www.mathworks.com>, The Math Works Inc., Natick, Massachusetts, USA, May 1999.
- [6] J.W. Eaton, "Octave: A High-Level Language for Numerical Computations," <ftp://ftp.che.wisc.edu/pub/octave>, University of Wisconsin-Madison, Madison, Wisconsin, USA, May 1999.
- [7] <http://ptolemy.eecs.berkeley.edu>, University of California Berkeley, Berkeley, California, USA, May 1999.
- [8] J. Ousterhout, "TCL Resource Center," <http://www.scriptics.com/resource>, Scriptics Corporation, Mountain View, California, USA, May 1999.
- [9] <http://www.java.sun.com>, Sun Microsystems, Inc., Palo Alto, California, USA, May 1999.
- [10] J. Shaffer, J. Hamaker and J. Picone, "Visualization of Signal Processing Concepts," *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 1853-1856, Seattle, Washington, USA, May 1998.
- [11] Y. Zeng, T. Weilundemo, R. Duncan and E. Wheeler, "Convolution Tool," <http://www.isip.msstate.edu/projects/speech/education/demos/convolution/index.html>, Mississippi State University, Mississippi State, Mississippi, USA, May 1999.
- [12] Y. Zeng, R. Brown and J. Hamaker, "System Response Tool," <http://www.isip.msstate.edu/projects/speech/education/demos/system/index.html>, Mississippi State University, Mississippi State, Mississippi, USA, May 1999.
- [13] E.S. Raymond, et al, "Open Source: the Future is Here," <http://www.opensource.org>, May 1999.

- [14] K. Shobaki, "Learn About and Experience Spoken Language Technology," <http://cslu.cse.ogi.edu/learn>, Oregon Graduate Institute of Science and Technology, Portland, Oregon, USA, May 1999.
- [15] R.A. Cole, et al, "Survey of the State of the Art in Human Language Technology," <http://cslu.cse.ogi.edu/HLTsurvey>, Oregon Graduate Institute of Science and Technology, Portland, Oregon, USA, May 1999.
- [16] "Signal Processing Information Base," <http://spib.rice.edu/spib.html>, Rice University, Houston, Texas, USA, May 1999.
- [17] D. Klatt, "Review of text-to-speech conversion for English," *Journal of the Acoustical Society of America*, vol. 82, pp. 737-793, September 1987.
- [18] R.A. Cole, "The Free Speech Journal," <http://cslu.cse.ogi.edu/fsj>, Oregon Graduate Institute of Science and Technology, Portland, Oregon, USA, May 1999.
- [19] H. Boulard, et al, "Speech Communication," <http://www.elsevier.nl/locate/specom>, Amsterdam, The Netherlands, May 1999.
- [20] J.P. Hosom, "Tutorials and Demos," <http://cslu.cse.ogi.edu/tutordemos>, Oregon Graduate Institute of Science and Technology, Portland, Oregon, USA, May 1999.
- [21] T. Carmell, et al, "Spectrogram Reading," http://cslu.cse.ogi.edu/tutordemos/SpectrogramReading/spectrogram_reading.html, Oregon Graduate Institute of Science and Technology, Portland, Oregon, USA, May 1999.
- [22] J.P. Hosom, et al, "Training Neural Networks for Speech Recognition," http://cslu.cse.ogi.edu/tutordemos/nnet_training/tutorial.html, Oregon Graduate Institute of Science and Technology, Portland, Oregon, USA, May 1999.
- [23] J. Picone, "Automatic Speech Recognition Resources," <http://www.isip.msstate.edu/projects/speech>, Mississippi State University, Mississippi State, Mississippi, USA, May 1999.
- [24] J. Hamaker and B. Brown, "Experiments," <http://www.isip.msstate.edu/projects/speech/experiments/index.html>, Mississippi State University, Mississippi State, Mississippi, USA, May 1999.
- [25] R. Duncan, "CVS Anonymous Download Instructions," http://www.isip.msstate.edu/projects/speech/support/info/cvs_instructions.html, Mississippi State University, Mississippi State, Mississippi, USA, May 1999.
- [26] D. Gibbon, "EAGLES NEST — EAGLET Term Database," <http://coral.lili.uni-bielefeld.de/~gibbon/EAGLES/SLWG/TERMBANK/interface.shtml>, University of Bielefeld, Bielefeld, Germany, May 1999.

- [27] T. DiGiovanni, "Virtual Solar System," <http://www.nationalgeographic.com/solarsystem/ax/high.html>, National Geographic Society, Washington, D.C., USA, May 1999.
- [28] "Do-it-yourself talking characters," <http://www.fluent-speech.com/demos.htm>, Fluent Speech Technologies, Beaverton, Oregon, USA, May 1999.
- [29] "Bringing the future of conversational computing to you today," <http://www.conversa.com/index.asp>, Conversational Computing Corporation, Redmond, Washington, USA, May 1999.
- [30] M. Kornbluh, J. Deller, J. Hansen, J. Grant, M. Seadle, C. Cieri, "A National Gallery of the Spoken Word," National Science Foundation Digital Libraries Program, July 1998.
- [31] T. DeFanti and D. Sandin, "Scientists in Wonderland: A Report on Visualization Applications in the CAVE Virtual Reality Environment," <http://www.evl.uic.edu/EVL/SHOWCASE/cave/vrpaper/report.html#wondermenu>, University of Illinois at Chicago, Chicago, Illinois, USA, May 1999.
- [32] C. Galica, "Glenn Learning Technologies Project," <http://www.grc.nasa.gov/WWW/K-12>, NASA John H. Glenn Research Center, Cleveland, Ohio, USA, May 1999.
- [33] J. Gradwohl, "Smithsonian Without Walls," <http://www.si.edu/organiza/museums/ripley/eap/start.htm>, S. Dillon Ripley Center, Washington DC, USA, May 1999.
- [34] B. Shelly and S. Bossert, "The Living Schoolbook," <http://lsb.syr.edu>, Syracuse University, Syracuse, New York, USA, May 1999.
- [35] R. McClintock and F. Moretti, "Institute for Learning Technologies," <http://www.ilt.columbia.edu>, Columbia University, New York, New York, USA, May 1999.
- [36] "Technology and Multimedia in Higher Educations," <http://www.uic.edu/depts/lib/arcade/edutech.html>, University of Illinois at Chicago, Chicago, Illinois, USA, May 1999.
- [37] "The NODE: Learning Technologies Network," <http://node.on.ca/tfl>, London, Ontario, Canada, May 1999.
- [38] "World Wide Web Course Tools," <http://www.webct.com/webct>, Vancouver, British Columbia, Canada, May 1999.
- [39] "Iowegian's dspGuru," <http://www.dspguru.com>, Iowegian International Corporation, Shawnee Mission, Kansas, USA, May 1999.
- [40] "Welcome to DSPNet," <http://www.dspnet.com>, TechOnLine, Inc., Waltham, Massachusetts, USA, May 1999.



Joseph Picone received his Ph.D. in Electrical Engineering from Illinois Institute of Technology in 1983. He is currently an Associate Professor in the Department of Electrical and Computer Engineering at Mississippi State University, where he also directs the Institute for Signal and Information Processing. He has previously been employed by Texas Instruments and AT&T Bell Laboratories. His primary research interest currently is the development of public domain speech recognition technology. Dr. Picone is a Senior Member of the IEEE and a registered Professional Engineer. He has also served in several capacities with the IEEE, including an associate editor of this journal. Dr. Picone has published more than 90 papers in the area of speech processing and has been awarded 8 patents.



Jonathan Hamaker is a Barrier Fellow and Ph.D. student in the Department of Electrical and Computer Engineering at Mississippi State University. He received his B.S. and M.S. from Mississippi State University in 1997 and 1999 respectively. He leads the Internet-Accessible Speech Technology Group which focuses on developing public-domain DSP tools and technology which are available via web interfaces. Mr. Hamaker's research interests lie in the development of robust acoustic models for large vocabulary conversational speech recognition. He has also recently completed an internship with Microsoft Corp. where he worked on the development of noise-robust solutions for speech recognition. Mr. Hamaker is a member of Eta Kappa Nu and is a student member of the IEEE.



Robert M. Brown, Jr. was born in 1974. He received a B.S. in electrical engineering, with an emphasis on communications, from Mississippi State University in 1998. From Sept. 1998 to April 1999, he worked for the Institute for Signal and Information Processing developing educational applets using the Java language. Since April 1999, he has worked at Digidesign, Inc. as a software engineer. His fields of interest are digital audio and acoustics. Previous research has included the development of linear discriminant analysis to help determine the scenic quality of forest images, and software to detect and remove acoustical feedback from a digital audio source. Mr. Brown is a member of the Institute for Electrical and Electronics Engineers and the Audio Engineering Society.



Ron Cole is a Research Professor in the Institute of Cognitive Science at the University of Colorado at Boulder. He is Director of the Center for Spoken Language Understanding at the University of Colorado, Boulder and the Oregon Graduate Institute of Science and Technology. He received his B.A. in Psychology in 1967 from the University of Rochester, his Master of Arts from the University of California Riverside in Psychology in 1969 and his Doctorate from UC Riverside in Psychology in 1971. His primary research interest is the development of interactive language technologies for classroom learning. Professor Cole has been actively involved in the study of speech for the past 25 years, resulting in over 100 published articles in books, scientific journals and conference proceedings.



Dr. John H.L. Hansen is an Associate Professor at the University of Colorado at Boulder. He holds appointments in both the Departments of Speech, Language, and Hearing Sciences, and Electrical and Computer Engineering. He serves as Associate Director for The Center for Spoken Language Understanding (CSLU), and Director of research activities of the Robust Speech Processing Lab (RSPL) at CSLU. He received the Ph.D. and M.S. degrees in Electrical Engineering from Georgia Institute of Technology in 1988 and 1983, and the B.S.E.E. degree from Rutgers University in 1982. He is the author of more than 110 peer reviewed journal and conference papers, and is coauthor of the textbook "Discrete-Time Processing of Speech Signals," (Prentice-Hall, 1993; IEEE Press 1999).


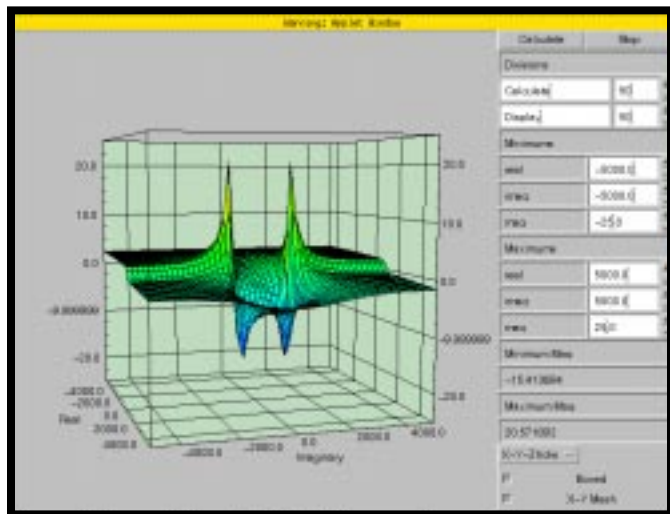
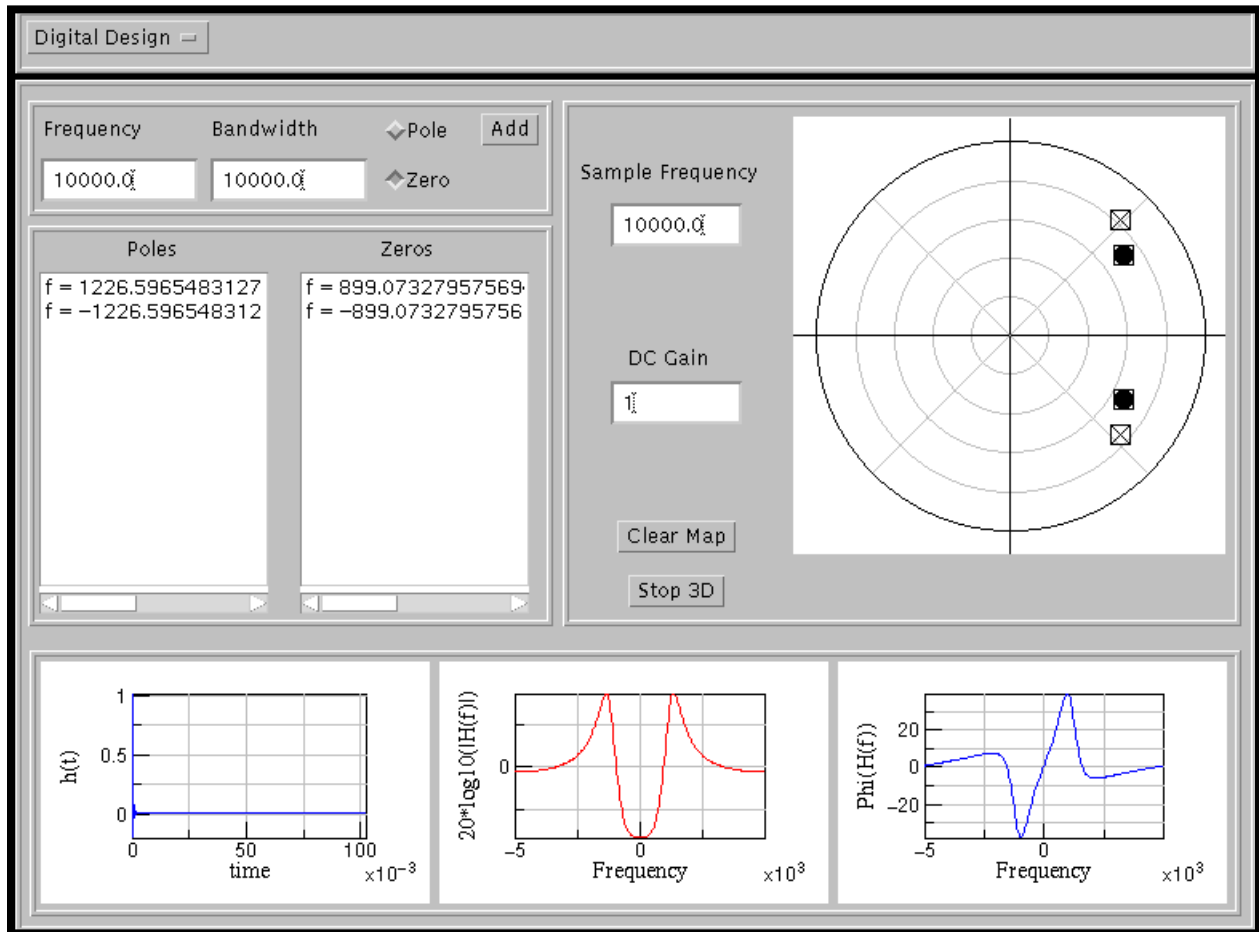


Figure 1. A more ambitious attempt to teach linear system theory using Java. In this applet, users can move poles and zeros around and see the impulse and frequency response of the resulting system change. Users can also interact with a 3D plot of the magnitude response.

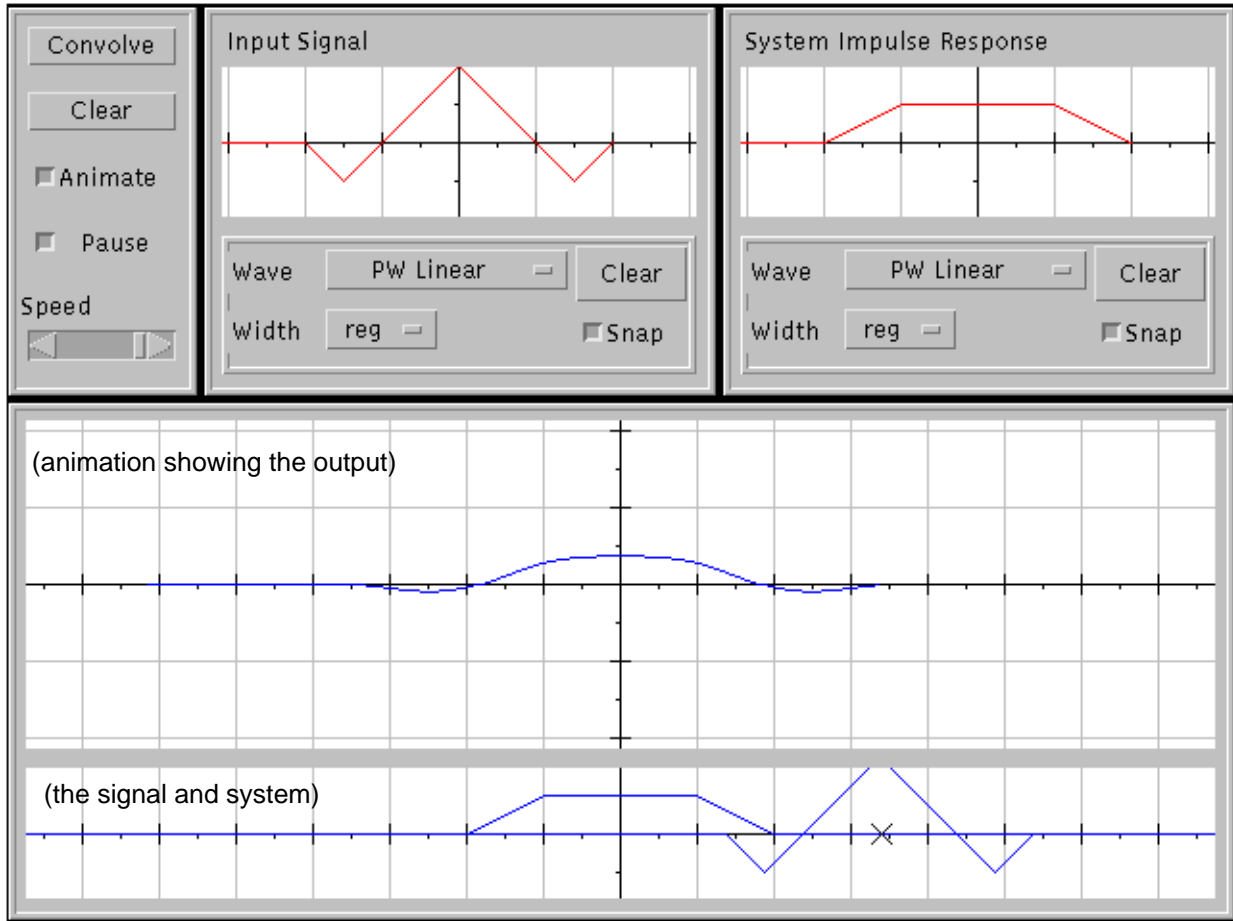


Figure 2. A Java-based interactive tool that teaches the concept of convolution.

The screenshot shows a Netscape browser window with the address bar displaying `http://cslu.cse.ogi.edu/tutordemos/rnet_training/tutorial.html`. The page content is titled "The Recognition Process" and features a flowchart with the following steps:

```

    graph TD
      A[digitize speech] --> B[compute spectral features]
      B --> C[classify time frames as phonetic categories]
      D[MLP] --> C
      C --> E[match category scores to target words]
      F[Viterbi search] --> E
      E --> G[measure confidence, accept or reject]
    
```

The flowchart is titled "The Recognition Process" and is attributed to the Oregon Graduate Institute of Science and Technology. The steps are: digitize speech, compute spectral features, classify time frames as phonetic categories (receiving input from an MLP), match category scores to target words (receiving input from Viterbi search), and finally measure confidence to accept or reject. The page also includes a sidebar with links to "Spectrogram Reading", "Robust Natural Language Parsing", "Neural-Network Recognition", "HMM and Neural-Network Recognition", "Text-to-Speech Demos", and "PSL Tools".

Figure 1. Diagram of the Recognition Process

There are four basic steps to performing recognition. Each step will be explained briefly here and in more detail later in this document.

First, we digitize the speech that we want to recognize; for telephone speech the sampling rate is 8000 samples per second. Second, we compute features that represent the spectral-domain content of the speech (regions of strong energy at particular frequencies). These features are computed every 10 msec, with one 10-msec section called a *frame*. Third, a neural network (also called an ANN, multi-layer perceptron, or MLP) is used to classify a set of these features into phonetic-based categories at each frame. Fourth, a Viterbi search is used to match the neural-network output scores to the target words (the words that are assumed to be in the input speech), in order to determine the word that was most likely uttered.

It is also possible to analyze the results by looking at the confidence we have in the top-scoring word. The word can be rejected as out-of-vocabulary if the confidence falls below a pre-determined threshold. The confidence values are always computed by the CSLU Toolkit during recognition, but we currently do not use these values during general-purpose recognition.

Figure 3. An example of an on-line tutorial on how to build a speech recognition system available as part of the CSLU toolkit [22].

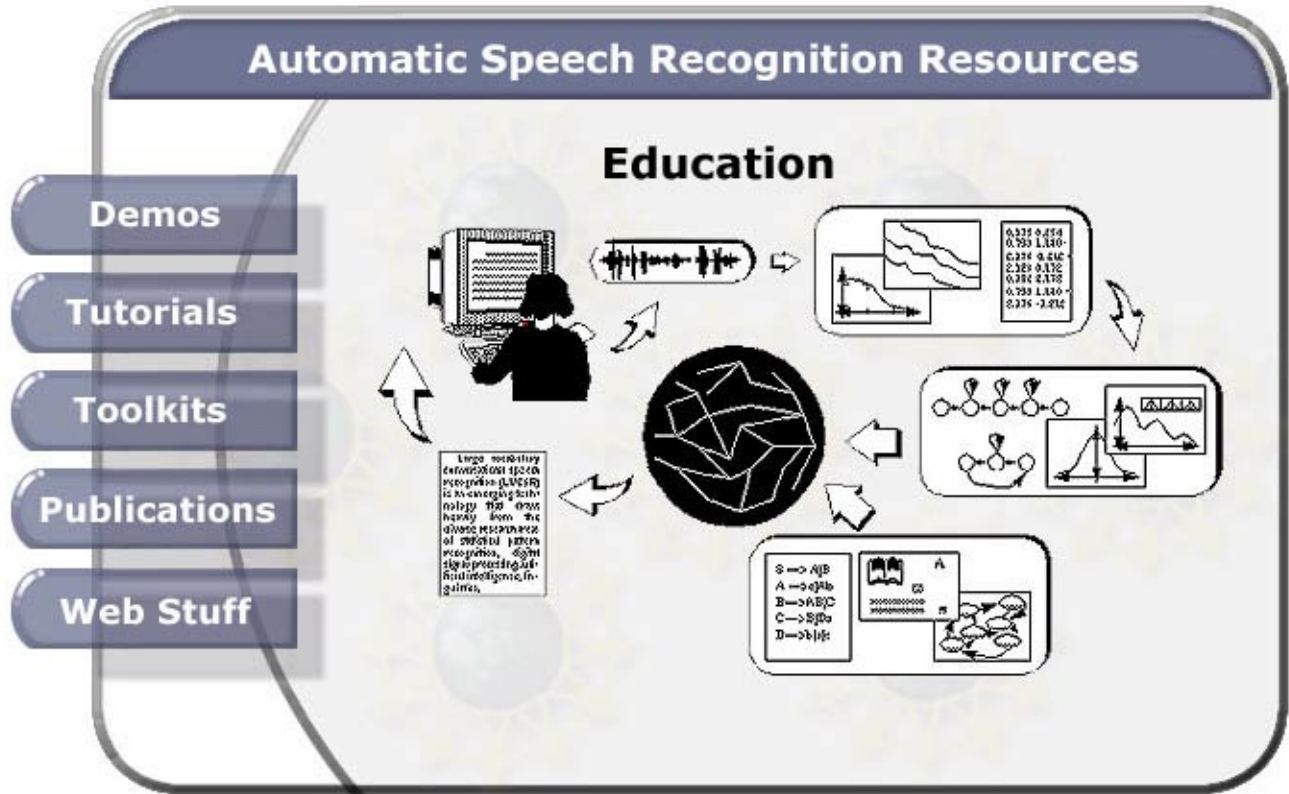


Figure 4. Successful web sites must offer a wide variety of information to its users in a compact fashion. This site devoted to speech recognition resources includes interactive demonstrations, on-line tutorials, toolkits to run specific applications of interest, and benchmarks that allow users to assess how new algorithms compare to existing state-of-the-art.

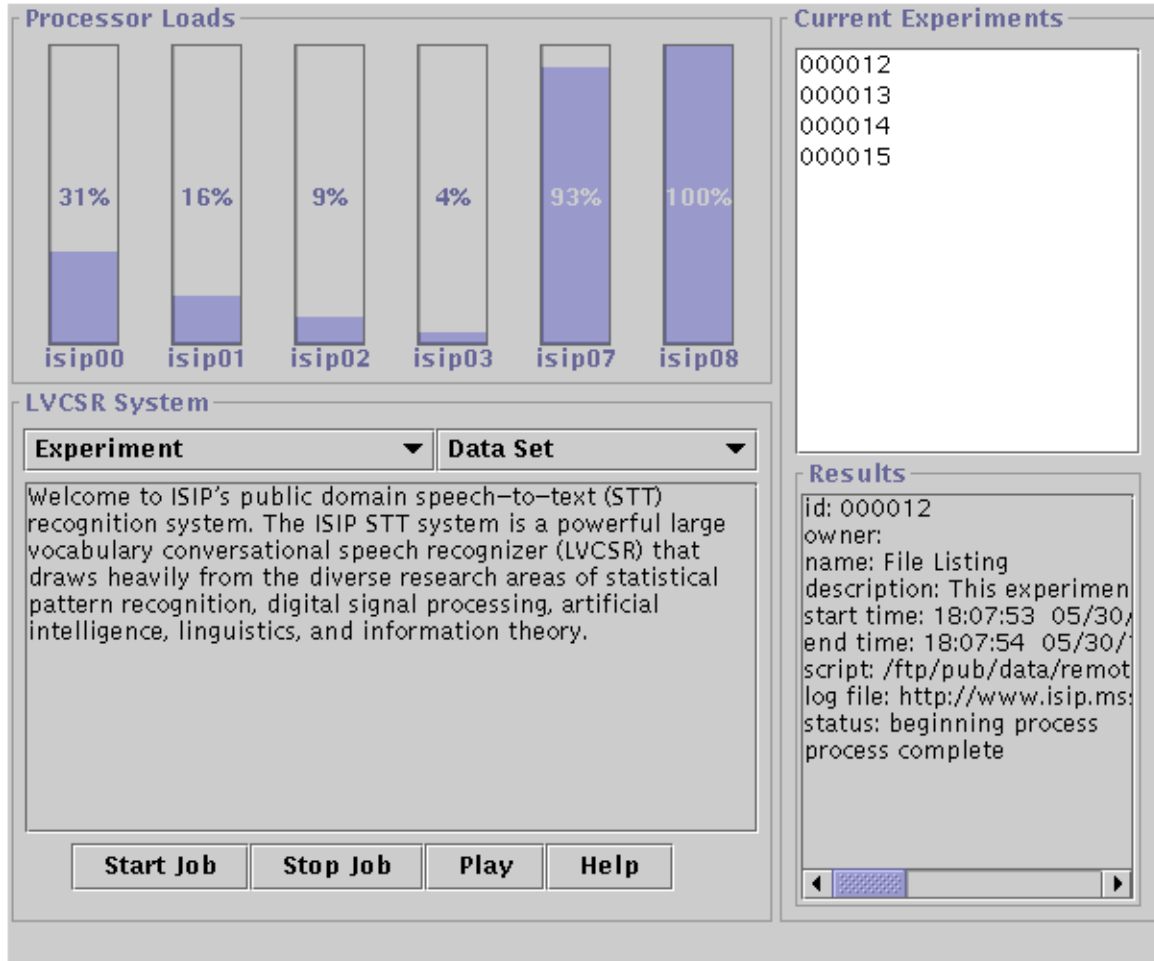


Figure 5. This web site, anticipating the growth in Internet bandwidth and capabilities, allows users to run speech recognition jobs over the Internet. Jobs using pre-stored data or a user's particular data can be submitted to a bank of servers. Results are returned via email and files that allow a user to analyze in great detail how the system processed the data.

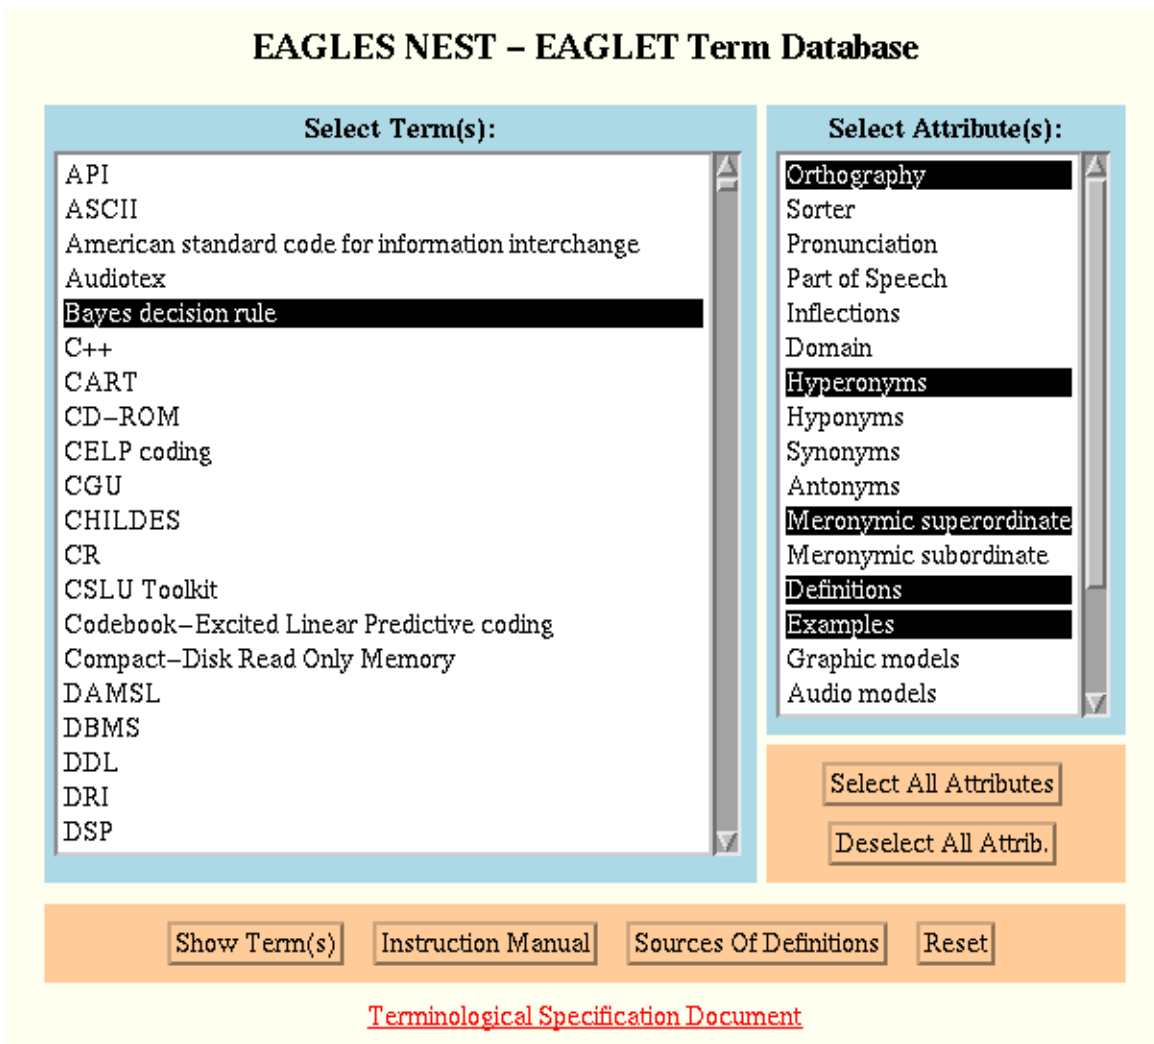


Figure 6. A terminology database that provides an extremely rich set of related links for a particular term. Such interfaces make it very easy to explore domains in a random manner.

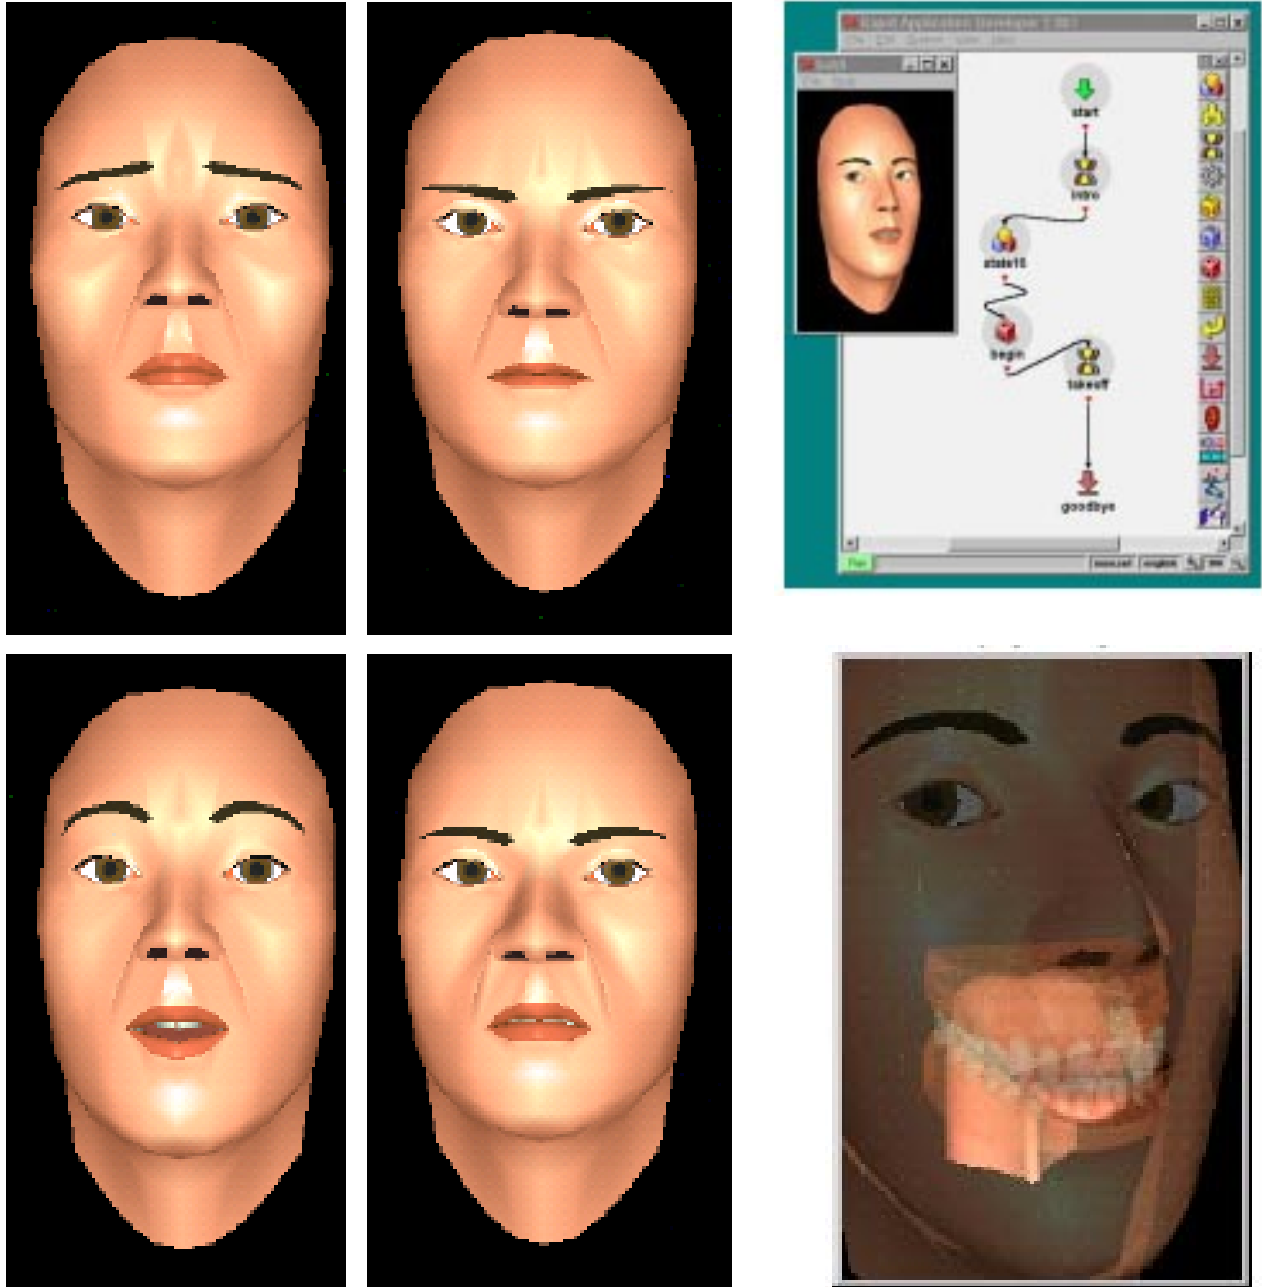


Figure 7. One example of “talking head” technology that is capable of expressing different emotions using control structures driven from models of paralinguistic information. Such interfaces add new dimensions to man-machine interaction, and have received enthusiastic response from users.




Figure 8. An example of a 3D immersive environment known as the CAVE. In such environments, users can visualize and interact with data with unprecedented realism. Significant cost barriers must obviously be overcome before such systems become accessible to the average educational institution.