

The World-Wide Web as a medium for psychoacoustical demonstrations and experiments: Experience and results

NORMA WELCH

McGill University, Montreal, Quebec, Canada

and

JOHN H. KRANTZ

Hanover College, Hanover, Indiana

The World-Wide Web provides an inviting opportunity to reach large numbers of people, both as an audience for psychoacoustical discussions and demonstrations and as subjects for psychoacoustical experiments. Such use of the Web poses a unique set of challenges. We present the methodologies used to prepare a multimedia primer in auditory perception at McGill University. As well as instructional material, this primer includes acoustical experiments conducted over the Web. We discuss the response and experiences with this site in general and in a particular class situation at Hanover College.

An increasingly popular addition to the Internet is the World-Wide Web (Web or the Web), which originated in March 1989 at CERN in Switzerland as a means of communication among high-energy particle physicists. Most information on the Web is in the form of hypermedia documents, written in hypertext markup language (HTML), and stored on a computer called a server. The reader accesses the information by running client software, called browsers, on any computer appropriately connected to the Internet. Popular browsers are Mosaic (NCSA, 1994) and Netscape (Netscape Navigator, 1994). These are available for many computer platforms.

Hypertext is text that contains user-addressable areas (mouse, hot keys, touch screens, etc.) which provide links to other places in the text or to other documents. Referenced materials may reside locally (on hard disk or CD-ROM), or remotely on any computer attached to the Internet. Hypermedia extends this concept by allowing links to other media, including graphics, playable audio files and video.

Central to the Web methodology is the standardization of the HTML protocols, through which client and server software communicate. The first standard, HTML 1.0, provided for hypermedia retrieval of information by the client. Only in 1994, with HTML 2.0, did communication from the client back to the server become possible: this protocol allows the server to cause a standardized form to be displayed on the client machine. Limited two-way communication between client and server was now possible. Even more flexible connections between client

and server are possible with more recent releases of Netscape and Mosaic.

The auditory perception tutorial and experiments at McGill University exploit the two-way communication now possible on the Web to conduct psychological experiments and to obtain feedback from readers on the effectiveness of auditory demonstrations. So also do experiments being conducted at Hanover College to examine the perception of attractiveness.

DEVELOPMENT OF THE MCGILL SITE

Initiation and Development

In January 1995, a graduate student special project was initiated with two major goals: to explain the concepts of auditory perception to a wide audience, and to test the feasibility of conducting experiments in auditory perception over the World-Wide Web. The first author of this paper developed the site found at <http://www.music.mcgill.ca/~welch/auditory/Auditory.html>. She had hoped to be able to use links to basic materials elsewhere on the Web and to focus on auditory stream analysis (Bregman, 1990), Shepard tone phenomena (Shepard, 1964), and the tritone paradox (Deutsch, 1991).

Unfortunately, this author was unable to locate any suitable introductory materials on the Web at that time. Accordingly, the site evolved to include tutorials in several basic concepts in auditory perception, as well as some historical information. The vocabulary of the multidisciplinary field of auditory perception is quite complex: an extensive glossary clarifies this language. The site also includes an annotated bibliography and extensive references. In June 1995, the site was discovered by the Darmstadt Auditory Research Group, Zoological Institute Darmstadt, who generously offered to create a

Correspondence concerning this article should be addressed to N. Welch, 2112 Claremont #29, Montreal, PQ, Canada H3Z 2P8 (e-mail: 75447.3462@compuserve.com).

European mirror site that is now available at <http://www.th-darmstadt.de/biologie/auditory/auditory.html>.

The majority of the materials on the site were developed between February and June of 1995. During this period, no attempt was made to publicize the site. Nonetheless, 173 accesses to documents (hits) within the Auditory World (the inner layer of the site; see below) were logged. In July, the site was registered with a handful of search engines and the number of hits rose to almost 1,000. In early August, the site was registered with many search engines and “What’s New” lists. Since then, hits have risen to over 1,500 a month.

Design of the Site

The Auditory Perception site is designed in two layers. The outer layer contains introductory material and hyper-text links to the following materials: Introduction, Some Selected Moments from History, Basic Acoustics and Psychoacoustics, Enter our Auditory World (to learn about and experience: Pitch perception, Auditory scene analysis, Dimensions of pitch and Shepard scales, The Tritone Paradox), Appendices, and other links to Web sites.

The real substance of the presentations and the experiments (and all audio examples) are in the inner layer. The reader has access to this through the Enter our Auditory World page. Client machines differ in the hardware and software used to interpret (and play) sound files, and in the audio formats that they can handle. The major formats used to transmit audio data are the following: .AIFF, suitable for Macintosh and SGI machines; .WAV, suitable for most PC systems; and .AU, suitable for Sun and NeXT computers. Rather than offering a choice of audio format within each demonstration, the Enter our Auditory World gateway allows the server to capture the reader’s preferred audio format once and for all, and thereby to simplify the interface within the individual demonstrations. Help and advice on audio formats are also provided from this gateway. This gateway serves also to capture the reader’s e-mail address or other identification for experimental purposes, and aids statistical analyses.

Access to topics within the Auditory World is governed by the following menu: Pitch perception (Pitch of simple tones, Pitch of harmonic tones, Residue or virtual pitch), Auditory scene analysis (The Auditory Scene Analysis Problem, Sequential organization, Organization of simultaneous components), Octave-related complex tones (Octave complex tones, Dimensions of pitch and Shepard Tones, The ever-ascending (or -descending) Shepard scale, Risset glides, The Tritone Paradox), Quick access to our Experiments (Hearing out harmonics, The Trill Experiment, Matching Octave-related complex tones, Does the Shepard effect persist with high “fundamental”?).

Each topic leads to a separate Web page. The design of each page takes advantage of the power of hypertext to allow the reader to tailor the readings to his/her level of expertise and interest: some contain lower level links to easy explanatory material, others contain links to sophisticated material (on pitch theories, for example), and some contain both easier and deeper links.

The material is intended to be read (and heard) in sequence, at least by readers unfamiliar with the field. However, the Web reader is naturally free to sample the materials in any order—and these readers do!

Preparation of Audio Examples

All of the audio signals were generated by using Csound (Vercoe, 1986), the most widely used sound synthesis and processing programming language. In this language, networks of tone-generating and -processing modules are assembled into *instruments*, in an *orchestra* file. The instruments then play from a *score* file—a time-ordered list of performance instructions. In these demonstrations, the instruments are mostly very simple, including, for example, sine tones, specified harmonic tones, noise.

Many of the demonstrations are repetitive and therefore require a set of similar scores. Accordingly, programs are used to produce them: the programs are written in Perl, a powerful interpretive language (Wall & Schwartz, 1992). There is essentially one script for each demonstration and experiment. In addition to producing a set of scores, these scripts “perform” the scores in Csound, convert the resulting audio files to the three audio formats, and move such files to their target locations—all without manual intervention. This automation reduced the time needed to prepare (or modify) a demonstration from several days to an hour or so.

All audio examples were monophonic and synthesized at a 22,050 sample rate and 16 bits per sample. This retains a good dynamic range and good quantization noise ratios, while providing up to 10-kHz bandwidth, which is sufficient for the selected demonstrations. Other trials with 8-bit data proved to be far too noisy. We have some concerns that listeners using low-quality acoustical equipment (such as inexpensive SoundBlaster type cards with budget speakers) will not hear them in the appropriate manner.

From the Appendix on the Web site, readers who are curious can view the orchestra files, the code that generated the score files, and a sample score file for each of the demonstrations and experiments. A significant number of readers have indeed done so.

Delivery of Audio

Two methods are used to allow the reader to access audio examples: clicking on a labeled loudspeaker icon, and clicking in a hot area of a *map*. A map is a graphic file embedded into the server environment and is sensitive to the physical area on the screen where the reader clicks with the mouse, much as one might click on a geographic map for detail about a geographic area. Use of a map is exemplified by the map of the chromatic circle of tones in Figure 1: clicking on any of the note names automatically fetches and plays the associated audio file.

Once selected, all audio is delivered to the client machine in the audio format selected by the reader. Depending on the setup of the client machine, it may be played immediately, or saved in a file from which it can be played repeatedly or processed by any local program.




Figure 1. An image map used to access sound files.

Currently (and unfortunately), audio files must be transferred in their entirety before they are played. Every effort has been made to keep the examples as short as possible, but the transfer time causes frustration to readers who have slow connections to the Internet.

Gathering Statistics

Because of the site's design, data on downloadings are generated on two separate files. Transfers of HTML, AU, WAV, AIFF, GIF, and JPEG files are gathered automatically by the server. Each downloading is date stamped and identifies the client machine. Once the reader has entered the auditory world, all HTML and audio files are generated by a program on the server; this program creates personalized HTML documents from a set of master HTMLs. The program inserts the user-selected audio format and buries identification data in each document. Audio format and identification are passed between documents. The disadvantage of this method is that data on downloaded HTMLs are no longer generated automatically on the server. Rather, the generating program itself accumulates data on client requests, together with experimental data, in a special log file, *auditory.log*.

All the programs controlling the site are Perl scripts residing on the server. One of these programs can display, for the interested reader, all the programs (including itself). Another displays various views of the log file and various views of the experimental data.

Another useful source of information has been e-mail messages received by the author. Finally, pointers to the site, especially those from educational institutions, indicate that it is fulfilling its role as a source of instructional material.

EXPERIENCE WITH THE SITE

Usage Statistics and Their Interpretation

In mid-October 1995, the number of lines in the *auditory.log* file (within the Auditory World) hit 5,000. This seemed to be a good time to study all the access statis-

tics, in order to determine whether the site was fulfilling its goals. Much care is needed in the interpretation of usage statistics. First, usage data is based on the client machine, and client machines do not map one-to-one with readers: many readers may access the site from the same machine, and a single reader may access the pages from several machines. A reader may even store URL references to some of the HTML pages that are generated within the Auditory World on the reader's hotlist and view them later from a different machine; thus we have occurrences of readers who view some of these internal pages, without appearing to pass through the gateway to them.

Second, readers differ widely in their behavior. Some return to the main menu before selecting the next item; others read to the end of the document and move on through the links that are provided to the next logical document. Even when readers return to the menu, this action may not be recorded in the server statistics, because the document has been stored on the client computer for reuse (cached). The result is that using the ratio of hits on internal pages versus the hits on the menu page gives an unreliable estimate of the seriousness of the average reader.

Third, some accesses may not be logged because the reader was connected through a proxy server which may deliver the documents from its store of documents (cache). With these points in mind, we have manually investigated a proportion of the data and written programs to automate the processing of the data in what we believe to be the most appropriate way. Given the large number of hits on the site (about 18,000 from about 3,500 machines), we believe that the rounded figures shown in Table 1 give a fair picture of the site usage.

There is, of course, no "typical" reader. Given the stereotype of the fast-moving cyberspacer who "surfs" the Web (always ready to jump onto the next interesting link), we are somewhat surprised at the serious interest that these usage statistics seem to indicate. To be sure, there is a marked fall-off of usage at each point: 29% of those who find the Home Page go no further; 30% of those who reach the gateway to the Auditory World do not enter; the first section within the Auditory World has 50% more hits than do the other two sections together; only 25% of those who enter the World stay longer than 20 min or look at more than six documents and only about 7% of them participate in the experiments. Some of the drop-off (see Table 2) can be accounted for by the number of readers who either have slow connections to the Web or have no audio equipment. On the other hand, 1,400 readers benefited by reading more than 5 HTML pages. Some readers got really involved in the topic, with 88 of them accessing more than 20 pages. A significant number of visitors have returned to the site. While some readers go straight to the Auditory World, others have spent considerable time in the outer area—reading the basics of Psychoacoustics or the historical information, or dwelling in the Glossary and reference lists. Half the activity in this area was by readers who looked at more than one of these documents.

Table 1
General Readership Statistics

Readers	3,500
Downloaded HTML documents	18,000
Readers of more than the Home Page	2,500 (71% of readers)
Readers of more than 5 pages	536 (15% of readers)
Readers of more than 10 pages	147 (4% of readers)
Readers approaching Enter Our Auditory World	1,600
Readers who actually entered Auditory World	1,050
“Serious” readers	245 (23% of those who entered)
Audio files downloaded	4,432
Listeners who downloaded audio files	717
Listeners who listened to more than five sounds	150
HTML documents read from the Home Page (basic theory, history, and appendices)	4,300
Readers of HTML documents from Home Page	1,850 (almost 50% of those who found the Home Page)
Readers of more than one HTML from Home Page	500

Note—“Serious” was defined as more than six HTML documents and/or longer than 20 min in the Auditory World.

The Discussion Corner, where readers can leave comments and questions for other readers, has been a complete failure. It has been accessed by a mere 10 readers, 5 of whom have left comments. Readers seem to prefer to submit questions and comments directly to the author.

The site has been used as reading material by a class in multimedia at the University of Southern California. Many of the students have written e-mail expressing appreciation of the materials. As will be discussed in the following sections, it has also been used for psychological experiments by a class at Hanover College. The high number of readers who selected the .WAV format indicates that many are accessing the Web from PCs.

Data From Experiments

Currently, four experiments are contained in the Auditory Perception pages: a trill threshold experiment, a test of hearing out harmonics in a complex harmonic tone, an investigation of how people hear octave-related complex tones, and an investigation of whether or not the Shepard tone effect persists at high frequencies. The number of results returned is shown in Table 2. This response to experiments is disappointing, but we do have 77 experi-

mental results. That is, about 7% of the readers who listened to some audio also responded to the experiments.

In the hearing harmonics experiment, readers were asked to report how many harmonics they were able to hear in a complex harmonic tone. About half reported being able to hear the second harmonic, and a few heard higher harmonics.

Results from the trill experiment are confusing. More than half the respondents supplied their e-mail address but failed to enter any data. We need to review the user interface here; perhaps the instructions were not clear. Of those supplying data, most heard the trill condition only with the least separated pair of tones—as expected.

Octave-complex tones are ambiguous, and this was confirmed by the variety of responses as well as the certainty expressed by the respondents. That we tend to hear in the 100- to 400-Hz range, even when the lowest pitch is below that, was confirmed by the results. It is disappointing that few of the readers who listened to these tones went on to read the section on pitch theories.

E-Mail Response

The ability of the site to elicit communication with the author has been very positive, with over 30 readers giving personal feedback. They had a varied experience in the area, ranging from novices to experts, and most of their comments were congratulatory. Some pointed out minor errors, which have since been corrected. Obtaining this feedback from the auditory community is a notable advantage to World-Wide Web authors. Some reported difficulties with downloading audio and/or with the quality that they were able to hear through their audio equipment. We are pleased that several readers have added pointers to the McGill site.

Discussion

The McGill site has, by any standards, been very popular. In addition, it has managed to hold the attention of a large number of readers from around the world. The number of readers who stay for a significant length of

Table 2
Usage of Auditory World (June to October 1995)

Readers entering the World	1,050
Hits on HTML documents	3,104 (disregarding menu)
By section:	
Pure and harmonic tones	1,853
Auditory streaming	526
Octave-complex tones and Shepard tones	680
By audio format	
.AU	42% of listeners
.WAV	40% of listeners
.AIFF	18% of listeners
Number of experimental results	77
Hearing harmonics	22
Trill experiment	35
Octave-complex tones	20

time (or revisit) indicates that it is fulfilling a useful role as a tutorial in auditory perception. At present, download times and lack of audio equipment limit the number of readers who can benefit.

The response to experiments, though low, is significant. Together with the number of e-mail responses and the number of comments on the experiments, this shows considerable potential for two-way communication (over the Web) between author and readers. The Discussion Corner was a failure.

EXPERIENCES AT HANOVER COLLEGE

This past fall, a sensation and perception class at Hanover College used the octave-complex experiment from the site. This experiment formed one of several laboratory experiences included in the class. In fact, the use of this site shows one of the great advantages of the Web for educational purposes. Since a Web site can be accessed around the world, institutions can thereby expand into areas limited in the past by their equipment. In addition, this site allowed the second author, the instructor, to more thoroughly examine auditory phenomena in the classroom. Though auditory demonstrations from CDs would have been possible at Hanover College, auditory laboratories were beyond the department's capabilities. So the use of an experiment over the Web allowed us to overcome one of the weaknesses of the perception course.

However, several obstacles needed to be overcome to use this laboratory. This institution does not have many computers with sound capability, and the psychology department does not at this time have its own set of networked computers. Thus, the entire class had to run on the instructor's computer to collect the data; this could have been difficult with a large class. Fortunately, there were only 12 students, and a network laboratory for the department is under construction.

Getting the students to use Web pages was easy. The basic point-and-click interface of Web browsers seems easy for students to learn. The delays in downloading some of the sound files were somewhat frustrating, but the browsers can be configured to keep the sound files in the sound player for easy repetition. A significant feature of the auditory perception experiments is the ability for a class to collect data under its own group names and thus segregate its own data. Thus, the students were able to collect their own data and use it in their laboratory reports.

The student reaction was generally positive. The ambiguity of the data results was frustrating but useful for teaching the students about individual differences, variability in perception, and the complexity of pitch perception. Therefore, it was a useful pedagogical tool. In addition, the students enjoyed using the sound feature, since multimedia is generally novel for students at present.

SUMMARY AND OUTLOOK

Whereas Hanover College has Web pages devoted solely to experimental data collection (Krantz, Scher, & Ballard, 1995), the McGill site combines demonstrations, explanations, and experiments. Both show that conducting experiments over the Web is certainly feasible. Our joint experience shows that data collection seems to proceed more reliably and with a greater frequency of obtaining data when the pages are devoted solely to the experiments.

The psychoacoustics experiment was very successful as a laboratory for a class, since this Web site has the ability to share the actual collected data and not just analyzed results. This is a strong and unique feature of the Auditory Perception site.

Audience response to the tutorials in auditory perception has been overwhelmingly positive, and the tutorials have been of clear benefit to a significant number of readers. They have also been successful in eliciting feedback to the author. We believe that the Web provides an excellent medium for making such materials available to a wide audience.

REFERENCES

- BREGMAN, A. S. (1990). *Auditory scene analysis: The perceptual organization of sound*. Cambridge, MA: MIT Press.
- DEUTSCH, D. (1991). The tritone paradox. In M. R. Jones & S. Hollerman (Eds.), *Cognitive bases for musical communications* (pp. 115-138). Washington, DC: American Psychological Association.
- KRANTZ, J. H., SCHER, J., & BALLARD, J. (1995, November). *Performing psychological research on the World-Wide Web: Some preliminary data*. Paper presented at the Indiana Academy of Science Fall Meeting, Indianapolis.
- SHEPARD, R. N. (1964). Circularity in judgments of relative pitch. *Journal of the Acoustical Society of America*, *36*, 2346-2353.
- VERCOE, B. (1986). *Csound*. Cambridge, MA: MIT, Media Laboratory.
- WALL, L., & SCHWARTZ, R. L. (1992). *Programming Perl*. Sebastopol, CA: O'Reilly & Associates.

(Manuscript received December 4, 1995;
revision accepted for publication February 8, 1996.)