mitted, and received at rates comparable to those of normal speech.

Perhaps the most important advantage of all is that the use of a pocket computer as a convenient, inexpensive communication device introduces the deaf TDD user to the concept of an intelligent, computer-based communication system of almost unlimited scope and flexibility. Breaking the cost barrier should be regarded as only the first step in developing considerably more powerful communication tools for the deaf. Devices such as these could do much to help break down the communication barriers between deaf and hearing users of the telephone system.

REFERENCES

[1]P. A. Bellefleur, "TTY Communication: Its History and Future," *Volta Rev.* **78**, 107-112 (1976).

[2]H. Levitt and J. R. Nelson, "Experimental Communication Aids for the Deaf," *IEEE Trans. Audio. Electroacoust.* **AU-18**, 2-6 (1970).

[3]J. Hurvitz and R. Carmen, *Special Devices for Hard of Hearing, Deaf, and Deaf-Blind Persons*, Little, Brown, and Co., Boston (1981).

LIP-READER TRAINER

Robin L. Hight
Varian Electric Research Co., St. Louis, Mo.

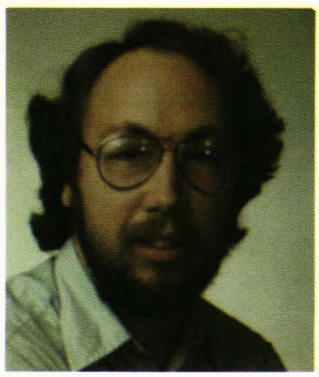

Robin L. Hight

Lip reading is the most difficult method of communication for the hearing-impaired person to learn. Once learned, however, it provides the hearing-impaired person with an important and valuable means of communicating with others. This paper presents an experimental software package that has been developed as a training aid to assist the parent or instructor in teaching lip reading to the hearing-impaired child or student. This article discusses the need for such a program, considerations of program design, a general description of the program, and future directions for the Lip-Reader Trainer.

Need for Assistance

Verbal communication between people is based on the transmission and reception of sound patterns that have been assigned definite meanings. These sound patterns are called words, and the sound segments that make up the words are called phonemes. When words are heard by a person, they are individually matched in memory to previously learned word patterns and assigned meanings. Lip reading requires a similar translation process for communicating except that, instead of matching auditory patterns with ideas or meanings, visual patterns are substituted. The visual patterns are formed during the utterances of the phonemes that create the words being spoken. Although there are over 60 phonemes in the English language, as few as 18 produce visually identifiable shapes. The rest either exhibit no visual features or take on the appearance of other phonemes. It is apparent that learning to associate phoneme shape patterns with words is quite difficult. Generally, an instructor or person knowledgeable in the correlation of phoneme shape to the phoneme is needed to teach the student. Further, much time is required to teach a person to lip-read, and an even greater amount of time is required for practice. Practicing also requires a willing and patient partner.

The Lip-Reader Trainer is an experimental software package for the Apple II microcomputer that converts typed sentences into animated mouth movements. It provides the instructor or parent with a training aid that presents the correct correlations of mouth configuration to phoneme, and easy-to-use animation controls. It also provides a flexible means to animate any words or sentences. In addition, the Lip-Reader Trainer is a patient partner who has as much time as the user wishes to take.

Design Criteria and Program Development

The design of the Lip-Reader Trainer had three main criteria: (a) to produce correct correlations of lip, teeth, and tongue configuration to phoneme; (b) to keep program memory requirements under 48K; and (c) to keep the programs easy to use. The first was the most difficult. Phoneme shapes are very personal in that everyone talks differently, even though they produce the same words. The position of the mouth, teeth, and tongue must form a certain configuration for each phoneme, yet the way that a phoneme appears can vary widely for different people. The phoneme shapes that were designed for the Lip-Reader Trainer are idealized and stylized shapes that clearly represent each of the 18 phonemes. Further, the shapes are presented from a straight-on view with the mouth at eye level.

The design of the shapes was a three-step process. First, *How to Read Lips* (Hawthorn Books, Inc., 1979) was referred to for a description of shapes for each of the 18 phonemes. Second, following those descriptions, the author sat in front of a mirror and reproduced the phoneme shapes. The shapes were then transposed to graph paper on a scale of four to one. The third step involved refining and defining the drawn shapes. The author was assisted during this phase of the development by an audiologist at the Central Institute for the Deaf in St. Louis. The main concern during the refining process was to make each shape easily identifiable without exaggeration. After three months, a set of phoneme shapes was finally developed for use in the Lip-Reader Trainer, including a nineteenth shape, which was neutral.

The second criterion for the design of the program was to keep the required memory storage within the stock capabilities of the Apple II computer. Memory space was needed for shape data, animation routines, the disk operating system, sentence conversion routines, and user programs, in addition to two open high-resolution graphic page areas. Storage of all the shape data in memory at one time was a very exacting requirement because shape access time had to be kept to a minimum in order for the animations to look smooth. Animation routines had to reside in memory in order to be able to sequence and draw the shapes. Sentence conversion routines also had to reside in memory to be able to control the animation routines. The disk operating system occupied its portion of memory and was needed to contain all the programs and sentences.

The first step in scheduling memory space was to convert the drawn phoneme shapes into shape data tables. Each shape was converted into graphic data by hand, assigned a storage location, and entered into memory. Each shape took about three hours to convert and store. Altogether, the 19 shapes occupied over 21,000 bytes of memory, or over half the usable memory space. Because we realized that the shapes used too much memory, a data compression technique was used to reduce the amount of data by half.

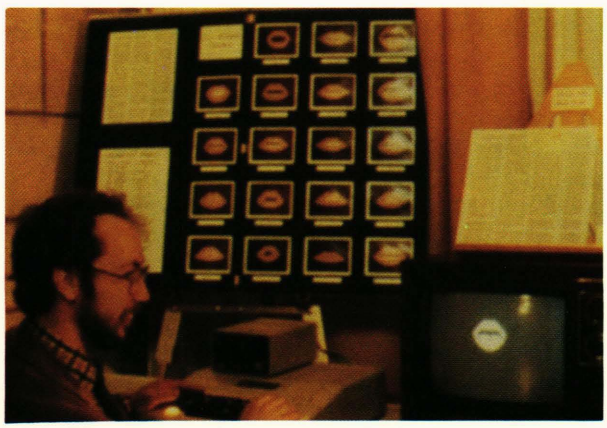

One of 19 phonemes having identifiable shapes is seen on the CRT of the Lip-Reader Trainer.

The second step in scheduling memory was to develop the necessary animation routine to reconstitute the condensed shape data. This program depends on a sentence conversion routine to tell it which shape to draw. Upon receiving instructions to draw a shape, the animation program finds the starting address of the requested shape in a "look-up" table. The shape is then reconstructed from the condensed data on one of the Apple II's high-resolution graphic pages. While this shape is being displayed to the viewer, the animation program receives instructions to draw another shape. This next shape is then drawn on the Apple II's other high-resolution graphic page and displayed. The animation program continues this process until it receives an "end of sentence" character, at which time the animation routine turns control back to the user program.

The third step was to develop a sentence conversion routine that would convert typed characters into the required phoneme shapes. The basis of this program is essentially a large look-up table. Each character in the sentence is equated to a specific shape. However, double letter sounds such as "ch" and "th" have to be handled slightly differently, as do diphthongs. As the sentence is converted, shape codes are generated and stored in a buffer area. It is from this buffer area that the animation routine receives its instructions as to which shape is to be drawn.

After these three programs were written and located in memory, it was determined that there was enough additional memory to support user programs. The user programs would bind together and make it easier to use the basic program components.

The third criterion was to develop easy-to-use programs for the user. Programs were developed for the instructor/parent and for the student. Instructor/parent-oriented programs consist of a Sentence Group Writer and an Evaluator. The student-oriented program is called Lip-Reader Trainer. The Sentence Group Writer provides the instructor/parent with a tool to create groups of words or sentences that can be stored on diskette and easily retrieved. The Lip-Reader Trainer provides the student with the means to retrieve the preprogrammed sentences and convert them into animated mouth movements. The Evaluator provides the instructor/parent with a record of the student's responses.

General Description

The Lip-Reader Trainer software package consists of three user programs and instructions written in BASIC, four system programs written in assembly language, 19 shape-data tables, and two directory tables. The program requires 48K of RAM, a disk-operating system, and two game-control paddles.

The Sentence Group Writer is used by the instructor/parent to create text files that will be used by the Lip-Reader Trainer as a source of sentences. For each sentence that is to be animated, the instructor must phonetically enter a master sentence, up to four

additional sentences, and a number. The additional sentences are used as multiple choice responses, and the number indicates which of the additional sentences is the same as the animated one. Several sentences can be entered to make up one sentence group in order to create a sense of context or focus on a particular problem. Further, by tailoring the response selection, the level of difficulty can be controlled for each master sentence. Since the software package is designed using the phonetic alphabet, it can produce animated sentences in many languages (if the other languages are written in phonetic English).

". . . the Lip-Reader Trainer is a patient partner who has as much time as the user wishes to take."

The Lip-Reader Trainer is the heart of this software package. It has the task of converting the phonetic sentences into animated mouth movements. This program also has the task of generating and maintaining student files. When the Trainer is accessed, it transfers the machine language subroutines and the memory from diskette to the Apple's memory. Next, the student enters his/her name and the sentence group number to be viewed. The sentence group number and the sentence number are displayed at the top of the screen. In the center of the screen the multiple choice responses are displayed, preceded by a message stating, "I am going to speak one of these." At the bottom of the screen the game control paddles are identified, and under each is a value from 0 to 255. The number represents a relative delay time between the displaying of each frame. One game paddle controls the length of time that each frame is displayed, and the other controls the delay between the last frame of one word and the first frame of the next word. In this manner, the student can control the speed of animation. Below the paddle setting display is a message for the student to press "return" when he or she is ready for the animation. When the return key is pressed, the screen is cleared and the program produces an animated mouth that speaks the preprogrammed sentence. When the animation is complete, the multiple choice sentences are shown and a response from the student is solicited. If the wrong answer is selected, the student is given the option to see the animation again, see the next sentence, or quit. If the first option is selected, the whole routine is repeated. If the second or third options or the correct answer are selected, entries are made in the student's file concerning the student's responses, and then the action selected is continued. These entries include date, student's name, sentence group number, sentence number; number of times the sentence was seen, status (CORRECT, GAVE UP, QUIT), letter speed, and word speed.

The Lip-Reader Trainer can also animate sentences entered directly from the keyboard. In this mode, the sentences must be entered in phonetics, and no multiple choice sentences are displayed. This feature can be used for impromptu sessions, for working out words that are not displayed correctly from dictionary phonetics, or for working out dialect displays. Sentences entered in this mode are not saved on diskette; however, the student file is updated.

The Evaluator is used by the instructor to retrieve the student data stored on diskette during practice sessions. The student's file is accessed by name and, optionally, by date. This allows the instructor to retrieve the most current records without having to page through the previous ones. The data presented, in conjunction with the instructor's notes on the sentence groups, allow the instructor to evaluate the student's progress or difficulties. Some general factors indicating student progress are how many times the sentence was seen, the status, and the relative delay times. Specific factors would include the level of difficulty that was programmed into the response selection or special sentences designed to drill particular phoneme patterns.

Future Development

Future development of the Lip-Reader Trainer will be made with the help of feedback from the hearing-impaired community. Feedback already received includes suggestions for full-face graphics, different viewing perspectives, a sentence library, speech synthesis (voice output), a sentence editor, and a printout of student files and sentence groups. Some developmental feedback may not be implemented because of restrictions inherent in the computer system (memory size).

The Lip-Reader Trainer has been modified to support a Votrax SC-01 type synthesizer. This modification requires that additional hardware be added to the Apple II in order to produce voice output. However, should a synthesizer not be available, the program will operate normally. Also, a sentence library has been added. The library is created from master sentences as they are added to sentence groups. Each sentence in the library is numbered and can be accessed for viewing, without updating a student's file, by specifying the sentence number.

Conclusion

It is not the purpose of the Lip-Reader Trainer to replace the instructor or to replace practice with people; it is meant to be an aid for teaching and a tool for practicing. Further, it is a program package that can be tailored and developed to help meet the needs of the hearing-impaired community.