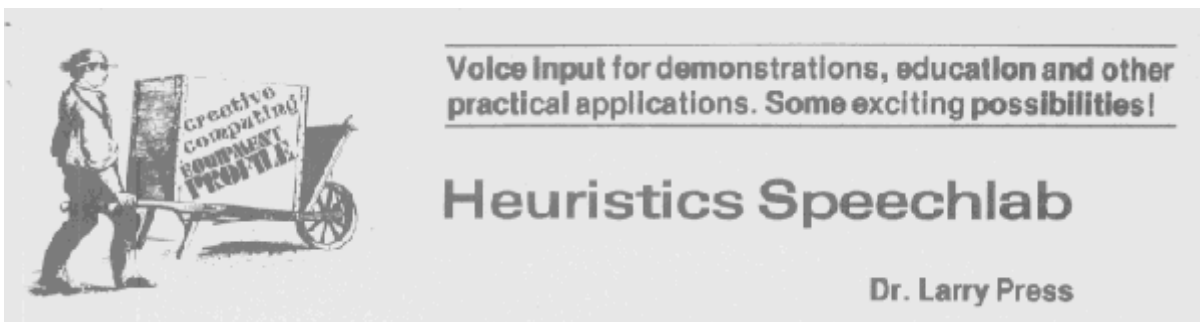


sp recog Heuristics

Heuristics Speechlab

dropbox



Speechlab,

a product of Heuristics Inc., is a voice recognition system (on a single S-100 board). It is used to digitize sound and pass the information to a computer which is programmed to analyze the sound data.

What Is It?

Figure 1 shows the relationship between the Speechlab System and an S-100 computer. The Computer sends one byte commands to Speechlab (" Beep your speaker ", " Read out a sample, " etc.) and receives digitized information on what is being spoken into the microphone. The interchange between Speechlab and the computer is accomplished via an I/O port which is on the Speechlab board.

Each time Speechlab is commanded to input information, it sends four bytes to the computer. One is an approximation of the overall frequency of the utterance since it was last queried and the other three approximate the amount of energy in three frequency ranges (averaged over time). These three ranges roughly bracket the first three resonances of the human vocal tract. Thus, if you are interested in listening to 2 seconds of speech and wished to sample the incoming sound every 10 milliseconds, you would end up with 800 bytes of information in memory.

In addition to hardware for digitizing sound in the manner outlined above, Speechlab comes with some useful software. It includes a demonstration program theoretically capable of learning and recognizing 64 different utterances, a subroutine for inputting speech (the number of samples to be input and the time interval between samples are the parameters), and a version of Palo Alto Tiny Basic with the speech input subroutine included. Last but not least, Speechlab comes with two manuals - one on assembly, test and principles of operations and a book of experiments. The manuals and software are outstanding and I will return to them later, but first let us ask who would want a Speech lab?



Who Is It For?

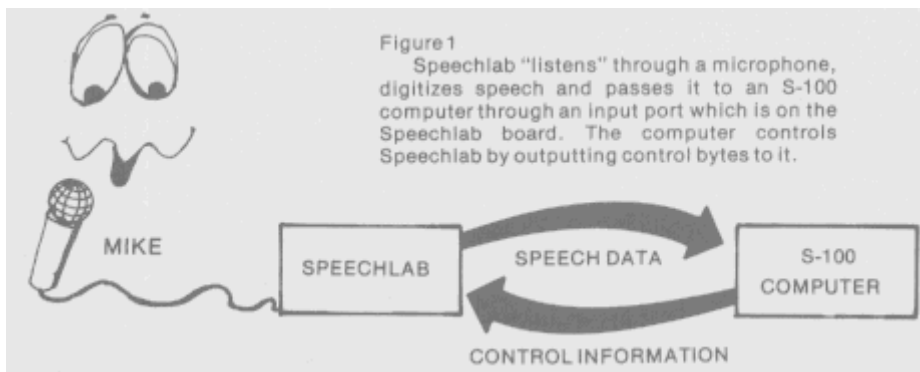
I can think of three reasons why a person might wish to purchase a Speech lab: for a cute demonstration, for some practical application, or in order to learn about linguistics and pattern recognition. I would rate Speechlab as "okay" for demonstrations, limited for practical applications, and fantastic as a learning tool.

If you are interested in a cute demonstration, you can have it using Speechlab and the demonstration program which comes with it. For instance, at a local computer store, they switch the fan on and off using Speechlab in conjunction with an AC controller (see photo). In a case such as this, where the system is trained to recognize one person's voice, where the ambient noise is low, where there are only a few utterances to distinguish amongst, and where the operator is able to alert the computer when it is time to "listen" for a command. Speech lab is fairly reliable and makes a satisfactory demonstration.

Unfortunately, if these conditions are not met, Speechlab is not very reliable. When I first tried the demonstration program which came with Speech lab, I was disappointed in that it became quite unreliable after training it to discriminate between only four or five phrases under ideal conditions. This limited performance could be due to either inherent limitations of the Speech lab system or to a poorly written demonstration program.

Inherent Limitations

In looking at the commercial speech recognition systems which are marketed by Threshold Technology and Scope Electronics for practical applications, we note that information is gathered in 16 and 16 frequency ranges respectively, rather than in just three as with Speechlab. This limited information is one inherent problem if you are thinking of practical applications



Dr. Larry Press, Box 5429, Santa Monica, CA. 90405
CREATIVE COMPUTING july 1979