Techniques for Eliminating Segmental Information

When performing perceptual experiments it is sometimes desirable to eliminate the segmental information from some part of the speech while preserving the amplitude envelope. Two techniques that have been used for this purpose are what we will call the Lehiste technique, in which the sign of every other sample is inverted, and the Schroeder/Samuel technique, in which the sign of random samples is inverted with probability 0.5.

The first set of figures are high-resolution sound pressure waveforms that resolve individual glottal pulses. The first is normal speech. The second shows the same speech after distortion by the Lehiste technique. The third shows the same speech as the first after distortion by the Schroeder/Samuel technique. We can see that both distortion techniques radically change the small-scale structure of the speech.

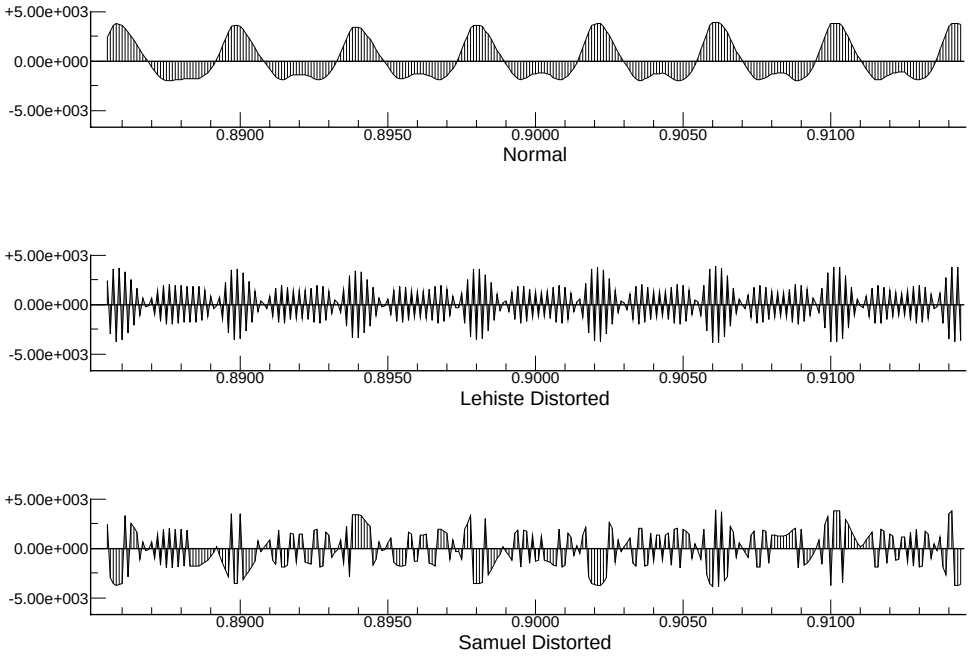

The next set of figures shows spectra computed over a window in the above speech samples. We see that the Lehiste technique inverts the spectrum, while the Schroeder-Samuel technique flattens it to white noise.

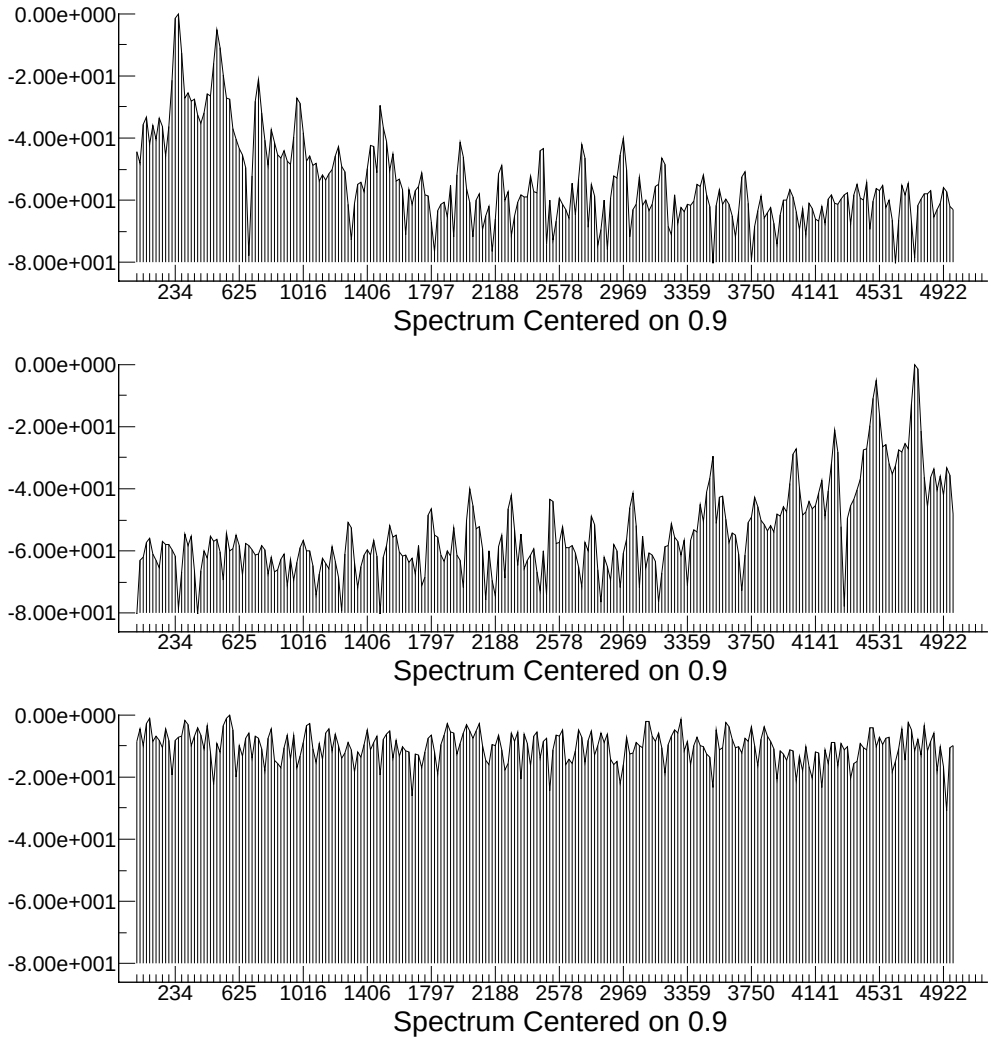

The next three figures give lower-resolution views of the same speech as above, showing that the amplitude envelopes are approximately the same.

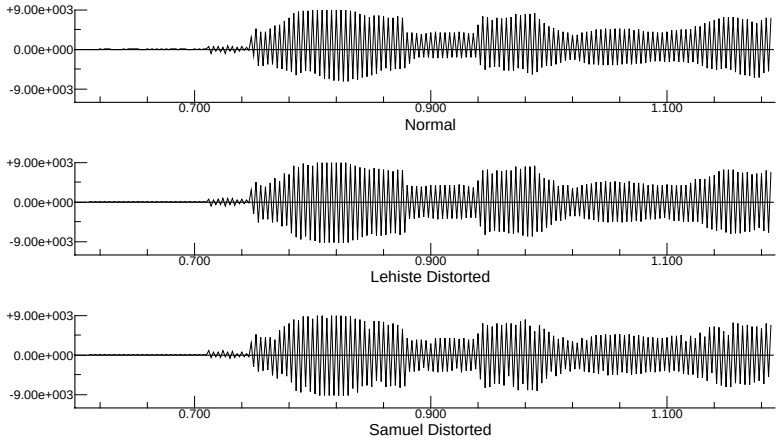

The Lehiste technique also preserves the F_0 contour, as can be seen by comparing the following two figures.

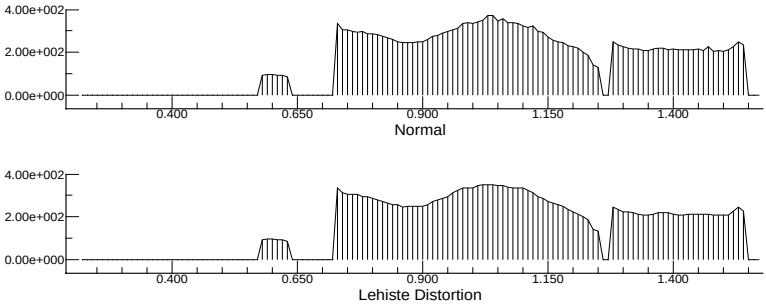

The Schroeder/Samuel technique does not preserve the F_0 contour since it produces a noise spectrum. The F_0 contour produced by the pitchtracker is essentially random. Here are the squared normalized autocorrelation coefficients associated with the successful period estimates produced by the autocorrelation pitch tracker that produced the above F_0 contours. These show that the original speech and the speech distorted by the Lehiste technique are generally fairly periodic, but that there is no detectable periodicity in the speech distorted by Samuel's technique.