

STQ AURORA DSR WORKING GROUP

Title: Experimental Framework for the Performance Evaluation of Speech Recognition Front-ends on a Large Vocabulary Task

Author: Günter Hirsch, Ericsson

Date: 4th June 2001

Version: 1.0

1 Summary

This report defines and describes the experiments for the recognition of a large vocabulary with the purpose of evaluating and comparing the robustness of feature extraction algorithms or complete recognition systems.

The WSJ0 (Wall Street Journal) database (available from LDC under the name CSR-I (WSJ0) complete) is chosen as basis for the experiments. The recognition of a 5000 word vocabulary is selected as task as it has been also used for the ARPA evaluations on continuous speech recognition. Besides the original data sampled at 16 kHz a second version of the data is created by downsampling the data from 16 kHz to 8 kHz. The reason for this is the required ability of the front-end to process speech at both sampling rates. Most of today's telecommunication terminals operate in the frequency range up to 4 kHz. But future speech services will aim at a higher speech intelligibility and a higher subjective speech quality by analysing speech at a higher bandwidth up to 8 kHz.

To evaluate the robustness in the presence of background noise different noises are artificially added to the original data. These noises represent realistic scenarios of application environments for mobile telephones. Noises have to be also available at 8 and 16 kHz sampling rate. The methods are described for estimating the energy of speech and noise signals and adding both signals at a desired SNR.

Details are presented about all recognition experiments. As already used in earlier experiments we define two training modes. One takes clean data only to train the recognizer. In the second mode clean as well as noisy data are applied to perform a multi-condition training. The predefined ARPA test set is also selected here to perform the recognition on a 5000 word vocabulary. Further noisy versions of this test set are created by adding individual noises at a certain range of SNRs. Recognition results are separately determined for each version.

2 Noisy speech database

The WSJ0 database is taken as basis for all experiments. Besides the original data sampled at 16 kHz with a precision of 16 Bit a second version is created by down sampling the original data to 8 kHz.

2.1 Filtering

An additional filtering is applied to consider the realistic frequency characteristics of terminals and equipment in the telecommunication area. Two "standard" frequency characteristics are used which have been defined by the ITU. The abbreviations G.712 and P.341 have been introduced as reference to these filters. Their frequency responses are shown in figure 1.

The G.712 characteristic is defined for the frequency range of the usual telephone bandwidth up to 4 kHz and has a flat characteristic in the range between 300 and 3400 Hz. P.341 is defined for the frequency range up to 8kHz and represents a band pass filter with a very low cut off frequency at the lower end and a cut off frequency at about 7 kHz at the higher end of the bandpass.




Figure 1: Frequency responses of G.712 and P341 filter

2.2 Noise Adding

Noise is artificially added to the clean WSJ0 data that have been recorded at a high SNR. To add noises at a desired SNR (signal-to-noise ratio) we define SNR as the ratio of signal to noise energy after filtering both signals with the G.712 filter characteristic. This can be applied to data sampled at 8 or 16 kHz. In the practical realization we use the available ITU tools for downsampling the data from 16 to 8 kHz and for applying the G712 filtering to data sampled at 8 kHz.

The speech energy is determined on basis of the ITU recommendation P.56 also available as part of the ITU software package. The noise energy is calculated as RMS value with the same software where a noise segment of same length than the speech signal is randomly cut out of the whole noise recording. We assume duration of the noise signal much longer than that of the speech signal.

The level of the speech signal is not changed as long as no overflow occurs in the Short-integer range. Based on the desired SNR the attenuation factor is calculated to multiply the noise samples before adding them to the speech samples. The speech level is only changed in case of an overflow. This happened only at a SNR of -5dB when creating the noisy versions of the TIDigits.

Noise signals are selected to represent probable application scenarios for telecommunication terminals. Noises have been chosen for the following environments or conditions:

- Car
- Crowd of people (babble)
- Restaurant
- Street
- Airport
- Train station

Some noises are fairly stationary like e.g. the car noise. Others contain non-stationary segments like e.g. the recordings on the street and at the airport.

Recordings in these environments are available e.g. from the NTT noise database. The NTT noises have been sampled at 22.05 kHz so that versions can be easily created at 8 and 16 kHz sampling rate. The long-term spectra of the noise signals are shown in figure 2 for the 6 noise conditions. These power density spectra have been estimated with the Welch method.

To add noise in case of 16 kHz sample rate speech and noise are downsampled to 8 kHz and filtered with the G.712 characteristic first to determine the weighting factor for the noise. Then speech and noise sampled at 16 kHz are filtered with the P.341 characteristic before adding them at the desired SNR.

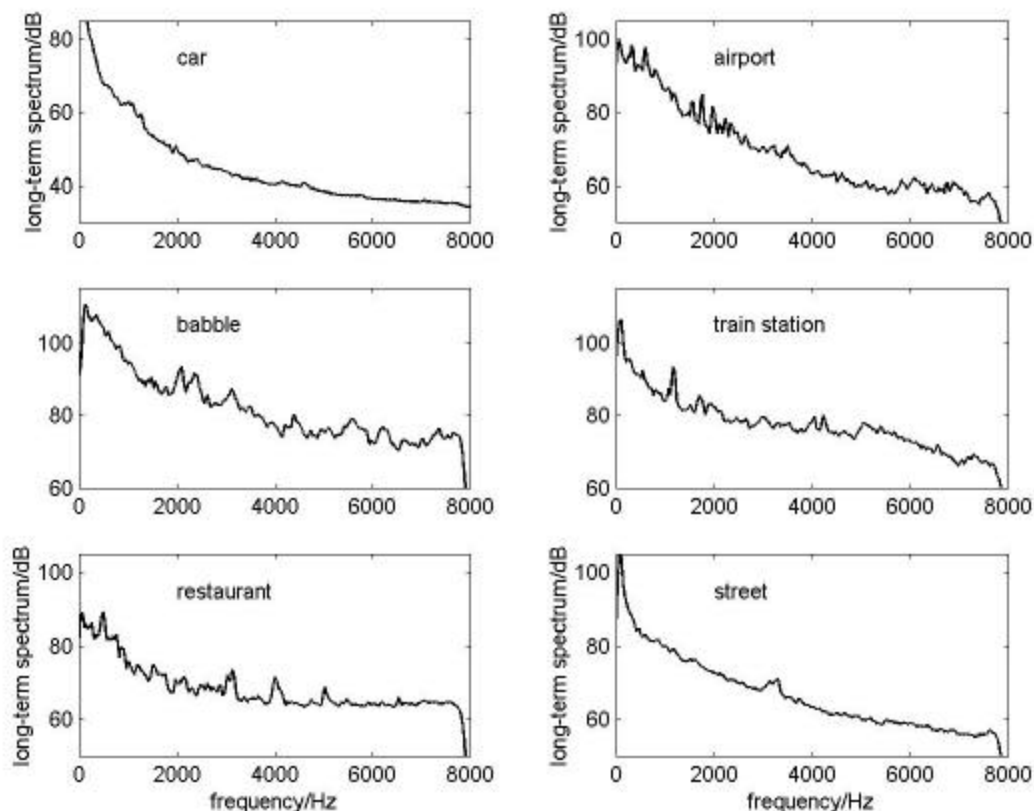


Figure 2: Estimated power density spectra for the 6 noise signals

3 Definition of training and test sets

The WSJ0 database consists of speech that has been recorded with two microphones in parallel. The first one is a close-talking microphone of type Sennheiser HMD414. No regulation exists about the choice of the second microphone which has been e.g. a desk mounted microphone. Most recordings consist of read texts from the Wall Street Journal.

In the ARPA evaluations a set of about 7200 utterances (~12 hours of speech) has been selected for the training. 7138 recordings are available on the CDs. These data are taken from the recordings with the Sennheiser microphone. We consider the same set for our training mode on **clean** data only. To define the multi-condition training we take the same set of 7138 utterances with half of them recorded with the Sennheiser microphone and the other half recorded with the second microphone. A variety of different microphones has been used as second microphone (e.g. Sony ECM-55, Sony ECM-50PS, Crown PZM-6FS, Crown PCC-160, RadioShack omni-electret, Nakamichi CM100, AT&T 5400 cordless phone, Panasonic KXT2365 speaker phone, ...). No further noise is added to one fourth of each of the two subsets. To the rest of the data noise is artificially added. The type of noise is randomly chosen out of 6 noise s in total and at a randomly chosen SNR between 10 and 20 dB. Goal

is an equal distribution of noise types and of SNRs over the whole range. Thus the average SNR is 15 dB. Figure 3 gives an overview about the data sets for the two training modes.

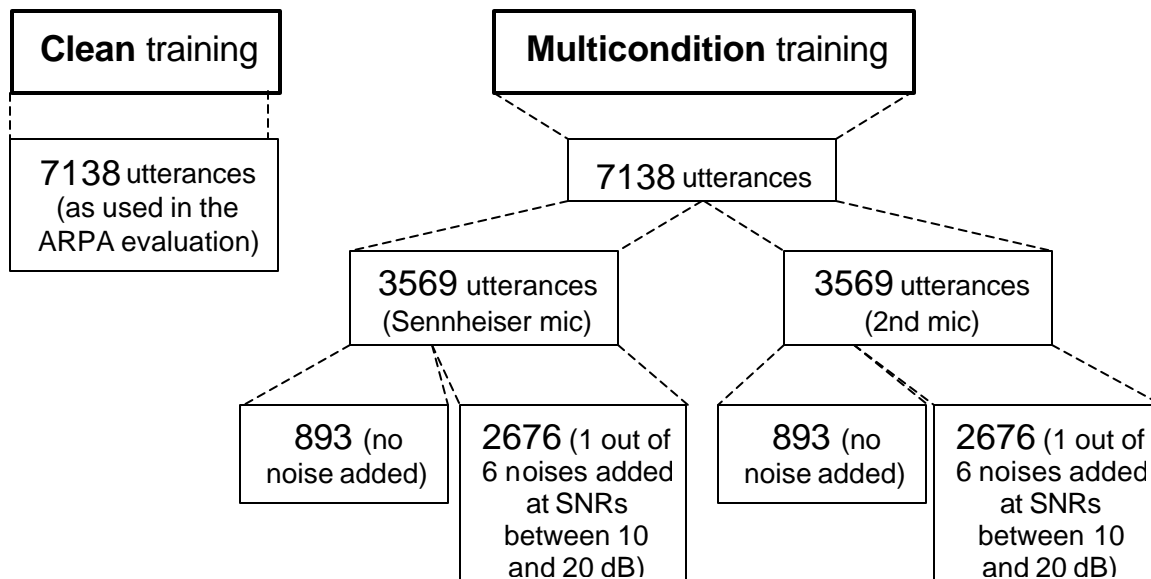


Figure 3: Data sets for the two training modes

A set of 330 utterances has been designated in the ARPA evaluation to perform a baseline recognition on the 5000 word vocabulary. This test includes the usage of a closed vocabulary bigram language model as supplied by Lincoln. The 330 utterances contain recordings from 8 speakers with about 40 utterances per speaker.

An alternative experiment has been proposed for the ARPA evaluation by using the same set of 330 utterances but recorded with the second microphone. Three different microphones have been used to record the 8 speakers in individual sessions. The same microphone has been used throughout each session. The frequency characteristics of these microphones have been analyzed in a differential way by comparing the parallel recordings with the Sennheiser microphone and the second microphone. A long term spectral analysis has been applied to the two versions of the same utterance. Comparing the two long-term spectra results in a rough estimate of the differential frequency response for each microphone in comparison to the Sennheiser microphone. The three differential frequency responses are plotted in figure 4.

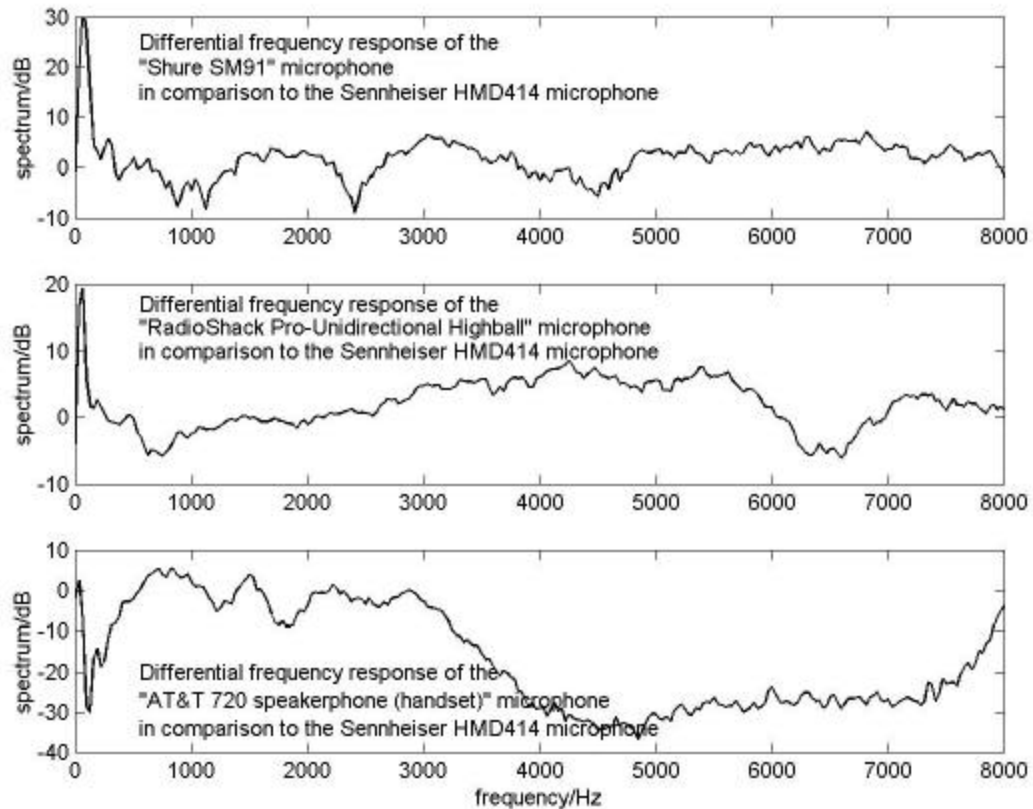


Figure 4: Differential frequency responses of different microphones

The peaks at very low frequencies should be ignored and are probably due to the highpass characteristics of both microphones. Having a big attenuation in this frequency region will cause an unreliable analysis result.

The highest influence on the spectrum has the AT&T handset microphone. It limits the speech signal to telephone bandwidth.

These two sets of speech data recorded either with the Sennheiser microphone or with one of the three other microphones are also selected for our experiments. Individual recognition results are created for each set.

6 further versions are created for each of the two sets so that we have 14 sets in total for testing. The 6 additional versions are created by artificially adding one type of noise for each version at a randomly chosen SNR between 5 and 15 dB. The goal is again an equal distribution of SNRs over the whole range so that the average SNR is 10 dB.

Thus we get 6 additional test sets for the speech data recorded with the Sennheiser microphone and 6 additional sets for the recordings with the second microphone. Each of these additional sets is characterized by being contaminated by one of the 6 noise types at an average SNR of 10 dB.

Figure 5 gives an overview about all 14 sets of test data.

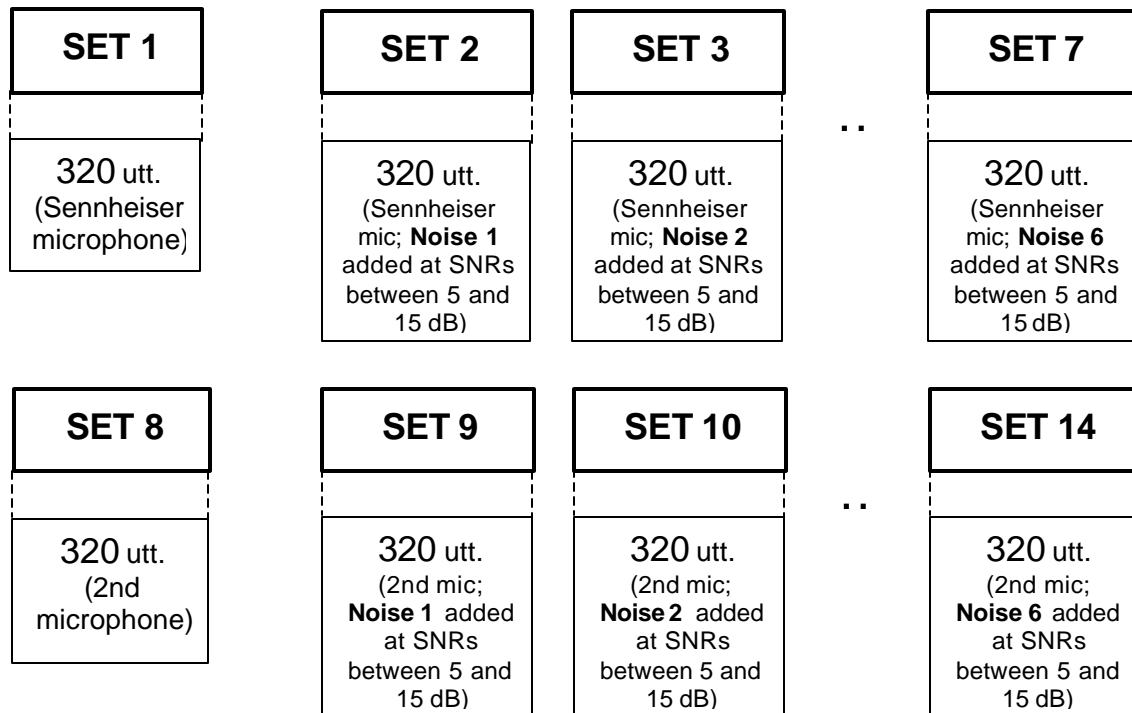


Figure 5: Data sets for recognition

As output we create 14 recognition results in total. The results for the two sets (sets 1 and 8 in figure 5) without noise added can also be taken to compare the performance against existing results of ARPA evaluations. All other results create a performance measure for individual noise scenarios as e.g. the car environment.

Considering the sampling at 8 and 16 kHz we have 4 sets of training data and 28 sets of test data.

We determine 28 recognition results at each sampling rate by recognizing the 14 test sets in both training modes.

4 ISIP recognizer

...