



Lombard Effect for Bilingual Speakers in Cantonese and English: importance of spectro-temporal features

Maximilian Karl Scharf¹, Sabine Hochmuth², Lena L.N. Wong³, Birger Kollmeier⁴, Anna Warzybok⁵

^{1,2,4,5}Medical physics and cluster of excellence Hearing4all, CvO University Oldenburg, Germany

³Clinical Hearing Sciences (CHearS) Laboratory, Faculty of Education, The University of Hong Kong, Hong Kong SAR, China

maximilian.scharf@uni-oldenburg.de, anna.warzybok-oetjen@uni-oldenburg.de

Abstract

For a better understanding of the mechanisms underlying speech perception and the contribution of different signal features, computational models of speech recognition have a long tradition in hearing research. Due to the diverse range of situations in which speech needs to be recognized, these models need to be generalizable across many acoustic conditions, speakers, and languages. This contribution examines the importance of different features for speech recognition predictions of plain and Lombard speech for English in comparison to Cantonese in stationary and modulated noise. While Cantonese is a tonal language that encodes information in spectro-temporal features, the Lombard effect is known to be associated with spectral changes in the speech signal. These contrasting properties of tonal languages and the Lombard effect form an interesting basis for the assessment of speech recognition models. Here, an automatic speech recognition-based (ASR) model using spectral or spectro-temporal features is evaluated with empirical data. The results indicate that spectro-temporal features are crucial in order to predict the speaker-specific speech recognition threshold SRT_{50} in both Cantonese and English as well as to account for the improvement of speech recognition in modulated noise, while effects due to Lombard speech can already be predicted by spectral features.

Index Terms: prediction of speech recognition, Lombard effect, tonal and non-tonal languages, speech features

1. Introduction

1.1. General aim

The aim of this study is to characterize the speech features involved in the Lombard effect for a comparison between a tonal language and a non-tonal language while minimizing effects due to variation between speakers.

1.2. Lombard Effect

Discovered in 1911 [1], the Lombard effect describes an increase in amplitude, duration as well as a change in the spectrum of speech, together with a change in visual cues when the speaker is introduced to loud sound fields. More precisely, the Lombard effect leads to a continuous upward shift of the fundamental F0 frequency and formants F1, F2 as well as an increased average vowel duration [2]. This change in speech production is also referred to as increased vocal effort. Additionally, it has been reported by [3] that the spectral change of increased vocal effort leads to improved speech recognition.

1.3. Tonal and non-tonal languages

Tonal languages, such as Cantonese, are spoken by a large fraction of the human world population and differ from non-tonal languages, such as English, in the property that information is encoded by changes in pitch within syllables. This has implications on the fundamental frequency f_0 which will be modulated over time to encode information. Thus, modeling a tonal encoding of information requires a model that incorporates both spectral as well as temporal features as an internal representation of speech signals. The contrasting property of tonality with spectro-temporal encoding of information and the Lombard effect, which acts on the long term spectrum of speech, makes it possible to assess the requirements on speech features that are used in speech recognition models. Established, and highly successful models for non-tonal languages are commonly used to describe speech recognition by spectral features only [4]. In the past, and with some success, these models have also been applied to tonal languages [5].

1.4. Spectral and spectro-temporal features of speech

The current literature (see above) suggests that an improvement of the speech recognition threshold due to Lombard speech is connected to spectral changes of speech, while tonal languages make use of spectro-temporal coding of information which remains mostly unchanged during Lombard speech.

Acoustic speech features are commonly derived from a short time Fourier transform (STFT) with logarithmically scaled frequency and amplitudes. For the frequency, an approximation of the perceptually motivated mel scale is typically used, which differs from \log_{10} in the base of the logarithm as well as in the vertical and horizontal intercept [6]:

$$\text{Mel}(f) = \log_{10}^{1/2595} \left(1 + \frac{f}{700} \right) \quad (1)$$

Automatic speech recognition systems can utilize this spectrum of a time-windowed section of a signal as feature vector, which is updated periodically for a fixed time shift and can be displayed on a 2D-grid as a spectrogram, called log-mel-spectrogram. A drawback of this data representation is that temporal changes do not enter the feature vector and need to be included by means of derivatives with respect to time.

One method to circumvent this problem is to calculate the Mel Frequency Cepstral Coefficients MFCC. MFCC take the log-mel-spectrogram and calculate the real part of its Fourier transform. Cepstra contain information about the temporal structure of a signal and allow to deconvolve, e.g., the glottal excitation signal from the vocal tract filter and the room impulse response that are characteristic for different frequencies. If more information about the temporal structure beyond the

range of the time window of the STFT is required, derivatives of the MFCC with respect to time are sometimes included in the feature vector.

Another method of collecting temporal information of speech for a feature vector is the use of Separable GaborFilter Banks, short SGBFB [7]. These work like a set of activation maps for the log-mel-spectrogram under the constraint that the resulting features are uncorrelated. SGBFB can be calculated efficiently, because of their property that the required 2D-filter can be separated into a filter for the time- and frequency- axis. Due to the fact that these activation maps spread over many time samples, it is expected that SGBFB contain more relevant temporal information for speech recognition than MFCC. SGBFB have been shown to be superior to MFCC for accurate simulations of speech recognition in fluctuating noise conditions [8].

1.5. Influence of noise on speech recognition

Besides the characteristics of the tested speech, the spectro-temporal structure of the masking noise is relevant for the resulting speech recognition threshold for which 50% of words are understood correctly (SRT_{50}). Generally, continuous noise leads to higher SRT_{50} than noise with a spectro-temporal structure, i.e., fluctuating noise, where prospective listeners are able to make extensive use of information that is present in sections of low SNR, which is also referred to as "listening in the gaps". For the purpose of this study, two types of a standardized noise were used [9, 10]: ICRA1, which is a stationary signal with speech like spectrum, and ICRA5-250, which has a temporal structure with gaps no longer than 250ms, such that the SRT_{50} of icra5-250 is expected to be lower than for ICRA1.

1.6. Hypothesis

This study aims to model the absolute SRT_{50} together with the Lombard Gain, which is the difference between SRT_{50} in Plain and Lombard speech, for bilingual speakers of Cantonese and English. Because of the different speech features that are relevant for the prediction of the Lombard gain and an absolute SRT_{50} prediction of tonal languages, the following hypotheses are tested:

1. spectro-temporal features are important to describe the speaker specific SRT_{50} across different types of masker while
2. spectral features suffice to describe the Lombard gain.

2. Methods

2.1. Empirical Data

Two female and two male bilingual talkers recorded the matrix sentence test corpus [11] in both Cantonese (Hong Kong variety YUE) and English for plain and Lombard speech each. For the purpose of this study, the Matrix sentence test was chosen, as it consists of a fixed grammatical structure with unpredictable content, which does not contain any information encoded in stress patterns. The Lombard effect was induced by ICRA1 noise with 80 dB SPL. For obtaining the SRT_{50} , these matrix sentence recognition tests were conducted with 15 normal hearing listeners, who were monolingual natives of the given language, in two different noise conditions: ICRA1 and ICRA5-250. The adaptive procedure with word-scoring used for this test was according to [12] and each adaptive track scanned manually for any abnormality prior to further processing. The absolute SRT_{50} of a speaker was calculated as the mean across all

listeners, while the Lombard gain was calculated as the mean across the individual Lombard gain for each listener. The measurement error of the mean σ_{μ} was estimated by:

$$\sigma_{\mu} \approx \sigma / \sqrt{N} \quad (2)$$

Where σ is the standard deviation across N listeners.

2.2. FADE Modeling approach

The Framework of Auditory Discrimination Experiments' (FADE) approach to speech recognition threshold modeling as proposed by [13] is based on an automatic speech recognizer (ASR), which acts as a model of a well-trained human listener. The FADE approach supports arbitrary types of features, based on the speech signals, to enter the ASR back end. For the simulation, the publicly available FADE distribution [14] was used with default settings.

The ASR-based SRT_{50} prediction consisted of the following steps: first, the audio of each sentence of the used test lists was mixed with random sections of the masking noise for N different SNR. Second, N Hidden Markov Models in conjunction with a Gaussian Mixture Model (HMM-GMM) were trained on the features that were derived from the labeled signals for each SNR. The log-mel-spectrogram, from which the features were calculated, was based on a STFT with 25ms hamming window and a time shift of 10ms. The underlying Markov model assumed 16 consecutive states per word and four states for silence, beginning or end of a sentence. Additionally, only transitions between words according to the matrix test syntax were possible. As a third step a new set of K matrix test audio signals, mixed with random sections of noise for a range of K different SNR, was used to test the performance of the N differently trained models. The resulting model performances as a function of N training- and K testing-SNR were used to calculate the SRT_{50} in a last step: the lowest value of the K testing SNR for a rate of 50% correctly identified words is returned as SRT_{50} . As the simulation is done for discrete values of testing SNR, the final prediction of the SRT_{50} is based on an interpolation of the data. Uncertainties on the model prediction originate from the size of the data set used for testing.

The FADE simulations were done for every combination of the two types of input features for the ASR (MFCC, SGBFB), speakers (4), language (ENG, YUE), vocal efforts (plain, Lombard), and masking noises (ICRA1, ICRA5-250), resulting in 64 computations.

2.3. Statistical measures

The predictive power of the simulations were quantified by the Pearson correlation coefficient R, the root mean square error RMS, bias of a linear regression and the χ^2/ν measure with ν degrees of freedom. χ^2/ν respects both measurement- σ_{emp} and simulation- uncertainty σ_{sim} :

$$\chi^2 = \sum_n \frac{(\text{sim}_n - \text{emp}_n)^2}{\sigma_{emp,n}^2 + \sigma_{sim,n}^2} \quad (3)$$

where sim_n denotes the result of a single simulation with index n and emp_n stands for the corresponding empirical measurement.

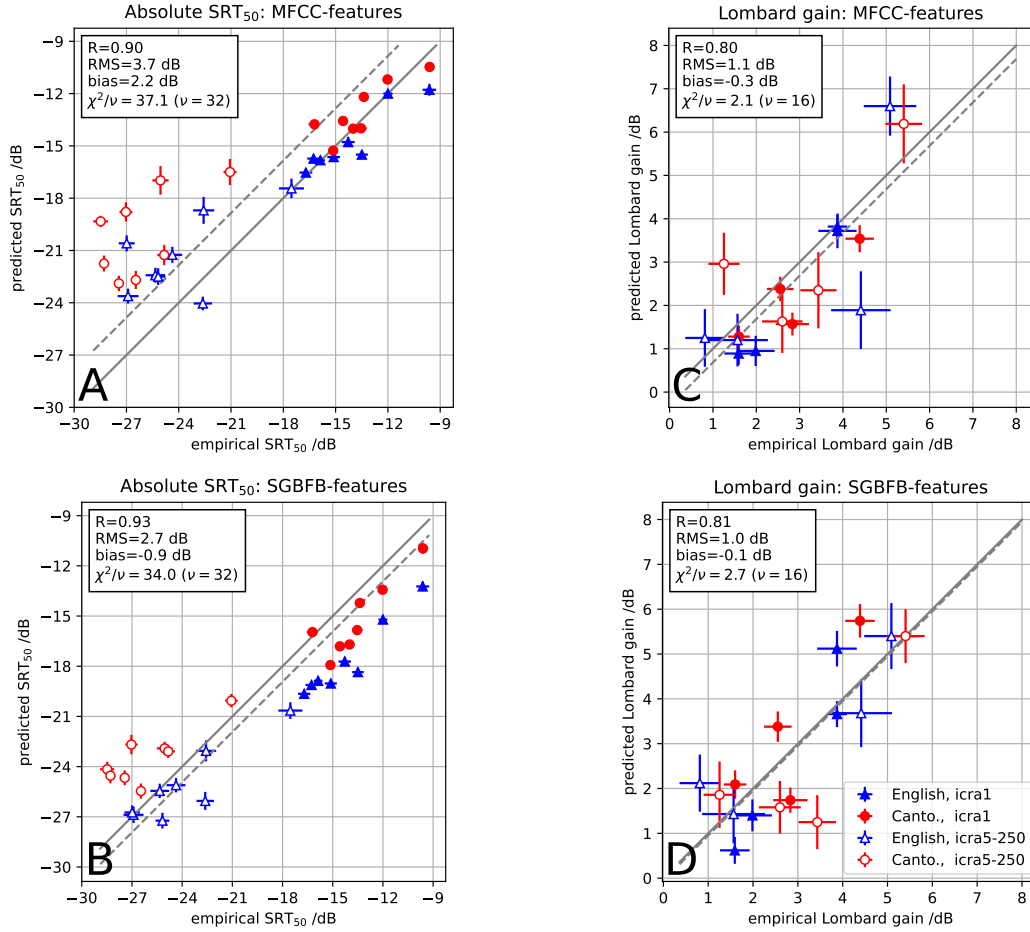


Figure 1: SRT_{50} and Lombard gain prediction of FADE with MFCC (A and C, respectively) and SGBFB features (B and D, respectively) for four bilingual speakers in stationary noise (ICRA1) and fluctuating noise (ICRA5-250). The solid line corresponds to perfect correlation, the dashed line depicts a linear interpolation.

3. Results

3.1. Simulation results

The scatter plots of predicted vs empirically obtained speech recognition thresholds (SRT_{50}) are shown in figure 1, where the prediction of the absolute SRT_{50} as well as Lombard Gain is displayed for MFCC- and SGBFB-based predictions. The predictions and empirical findings of absolute SRT_{50} correlate with $R \geq 0.9$ for both feature sets.

Panel A and B show that the FADE prediction based on SGBFB features was able to describe the data with smallest RMS, offset and χ^2/ν where ν is the degrees of freedom of the modelled data. As the FADE prediction method is only working with the sound files, no information about the empirical SRT_{50} has flown into the prediction, such that ν is equal to the number of data points.

The MFCC based prediction in panel A showed a good agreement between simulations and empirical data in stationary noise ICRA1 ($R=0.85$). However, the correlation is weak for the ICRA5-250 noise condition ($R=0.44$), which is line with previous findings [8] that FADE can predict the SRT_{50} in fluctuating noise better with spectro-temporal feature sets (e.g., SGBFB) than with either temporal or spectral feature sets in isolation.

The MFCC are not sufficient to take advantage of the temporal gaps in the masker as efficient as human listeners did. This observation holds for both languages. On the other hand, the predictions based on SGBFB features of panel B coincide very well with the empirical data for both noise conditions. A negative bias of -2.6 dB for ICRA1 remained for SGBFB features.

Panel C and D show the Lombard gain, which is the difference between the SRT_{50} in plain and Lombard speech. Accordingly the remaining degrees of freedom of the data ν are cut in half. Both feature sets show high correlation with vanishing bias together with a value of χ^2/ν close to unity, which means that the deviation of the model from perfect correlation with the measurement is in agreement with the uncertainties. It can be observed that the Lombard gain prediction was independent of the used features.

4. Discussion

The discussion will distinguish between the absolute SRT_{50} prediction and prediction of the Lombard gain according to the two initial hypotheses that spectro-temporal features are important to describe the speaker specific SRT_{50} while spectral features suffice to describe the Lombard gain.

4.1. Absolute SRT₅₀ prediction

Panel A and B of figure 1 show the absolute SRT₅₀ prediction for the two different speech features MFCC in A and SGBFB in B.

Based on the calculated statistical measures, the SGBFB based prediction was superior to the MFCC. Especially in the fluctuating noise condition ICRA5-250, MFCC based predictions failed to describe the empirical data, which is in agreement with known properties of MFCC [8]. However, SRT₅₀ in ICRA1 was correctly predicted with MFCC for both languages, which means that the MFCC features provided to the ASR backend of FADE were sufficient to predict SRT₅₀ for both English and the tonal language Cantonese in stationary ICRA1 noise. On the other hand, the SGBFB based prediction worked for all noise conditions, such that the individual SRT₅₀ of a speaker was predicted with higher accuracy, which is in line with hypothesis 1, that spectro-temporal features are important to describe the individual SRT₅₀. Even so, a bias of -2.6 dB is present for the SGBFB based prediction in the stationary ICRA1 noise, that is expected: FADE works as a model of an ideally trained listener, which may lead to a negative offset, if the listeners employed are neither highly trained nor selected to be free from any minor sign of a hearing deficiency. Note that according to [15], even normal-hearing listeners have to be modeled with a small processing deficit (denoted as level uncertainty) to obtain precise model predictions with FADE.

Even though SGBFB features resulted in the best prediction of the absolute SRT₅₀, as can be seen in figure 1 panel B, χ^2/ν is not in the range of unity, which shows that the model, or the applied features, still need improvement to describe the data to full agreement with measurement uncertainties.

4.2. Lombard gain prediction

Panel C and D of figure 1 show the Lombard gain prediction for the two different speech features MFCC in C and SGBFB in D.

As can be seen from the statistical measures with similar correlation coefficient, RMS of about 1 dB and bias below 0.3 dB, the two prediction methods are equivalent in their predictive power. χ^2/ν is close to unity for both models, which means that the deviation of the prediction from perfect correlation with the empirical data is in agreement with measurement uncertainties together with the precision of the model. It is evident that the Lombard gain, which is associated with spectral changes of speech, can be adequately modelled by MFCC. MFCC can only contain information from within the time window of the STFT, which was in this case 25ms. These findings provide evidence for hypothesis 2 that spectral features suffice to describe the Lombard gain.

5. Conclusions

The present study has shown that ASR based models of speech recognition such as FADE are a powerful tool for the prediction of SRT₅₀ in stationary and fluctuating noise conditions for Cantonese and English. However the choice of input features is important for a correct prediction of the individual SRT₅₀ of a speaker. SGBFB features, which are similar to activation maps for the log-mel-spectrogram, were superior to MFCC, which can only contain information limited by the time window of the STFT. On the other hand, it was shown that the Lombard gain can be predicted by both MFCC as well as SGBFB, because changes in SRT₅₀ due to the Lombard effect appear to be mainly associated to a spectral change of speech. In summary

it was found that:

1. spectro-temporal features such as SGBFB are important to predict the SRT₅₀ for both Cantonese and English for speakers in stationary and in modulated noise.
2. spectral features such as MFCC suffice to describe the Lombard gain of a speaker, for both stationary and modulated noise.

Although SGBFB have been shown to be superior to MFCC for individual SRT₅₀ prediction, the model was not able to compute predictions with maximum possible agreement with empirical data, which will require further investigations of the model or the used features.

Furthermore, the aspect of differences between tonal and non-tonal languages has not yet been fully explored in this contribution. Since bilingual speakers were used, further analysis of the data may uncover if significant differences in the Lombard Gain for Cantonese and English exist even for the same speaker. This might be related to the usage of different speech cues for enhancing speech recognition via the Lombard effect across languages.

6. Abbreviations

Abbreviation	Meaning
ASR	Automatic Speech Recognizer
FADE	Framework of Auditory Discrimination Experiments
HMM-GMM	Hidden Markov Model in conjunction with a Gaussian Mixture Model
ICRA1	unmodulated speech noise with a long term spectrum of a male speaker at normal vocal effort as defined by the International Collegium for Rehabilitation Audiology [9]
ICRA5-250	modulated speech noise with a long term spectrum of a male speaker at normal vocal effort as defined by the International Collegium for Rehabilitation Audiology [9] with a duration of pauses no longer than 250ms as proposed by [10]
MFCC	Mel Frequency Cepstral Coefficients
RMS	Root Mean Square error
SGBFB	Separable Gabor FilterBank
SNR	Signal to Noise Ratio /dB
SRT ₅₀	Speech Recognition Threshold for 50 % correctly understood words
STFT	Short Time Fourier Transform
YUE	a dialect of Cantonese that is spoken in Hong Kong

7. Funding

This work was supported by the DFG grant "Experiments and models of speech recognition across tonal and non-tonal language systems (EMSATON)" grant number 415895050.

8. Declaration of Conflicting Interests

The Authors declare that there is no conflict of interest.

9. References

- [1] E. Lombard, "Le signe de l'elevation de la voix," *Ann. Mal. de L'Oreille et du Larynx*, pp. 101–119, 1911.
- [2] N. Alghamdi, S. Maddock, R. Marxer, J. Barker, and G. J. Brown, "A corpus of audio-visual lombard speech with frontal and profile views," *The Journal of the Acoustical Society of America*, vol. 143, no. 6, pp. EL523–EL529, jun 2018.
- [3] J.-C. Junqua, "The lombard reflex and its role on human listeners and automatic speech recognizers," *The Journal of the Acoustical Society of America*, vol. 93, no. 1, pp. 510–524, jan 1993.
- [4] ANSIS3.5, *Methods for calculating the speech intelligibility index SII*, rev. 1986 ed., Acoustical Society of America, American National Standard, Jun. 1997.
- [5] J. Chen, Q. Huang, and X. Wu, "Frequency importance function of the speech intelligibility index for mandarin chinese," *Speech Communication*, vol. 83, pp. 94–103, oct 2016.
- [6] ETSI, "Speech processing, transmission and quality aspects (stp); distributed speech recognition; frontend feature extraction algorithm; compression algorithms," *ETSI ES*, vol. 201, no. 108, p. V1, 2003.
- [7] M. R. Schädler and B. Kollmeier, "Separable spectro-temporal gabor filter bank features: Reducing the complexity of robust features for automatic speech recognition," *The Journal of the Acoustical Society of America*, vol. 137, no. 4, pp. 2047–2059, apr 2015.
- [8] M. R. Schädler, A. Warzybok, S. D. Ewert, and B. Kollmeier, "A simulation framework for auditory discrimination experiments: Revealing the importance of across-frequency processing in speech perception," *The Journal of the Acoustical Society of America*, vol. 139, no. 5, pp. 2708–2722, may 2016.
- [9] W. A. Dreschler, H. Verschuure, C. Ludvigsen, and S. Westermann, "ICRA noises: Artificial noise signals with speech-like spectral and temporal properties for hearing instrument assessment: Ruidos ICRA: Señales de ruido artificial con espectro similar al habla y propiedades temporales para pruebas de instrumentos auditivos," *International Journal of Audiology*, vol. 40, no. 3, pp. 148–157, jan 2001.
- [10] K. C. Wagener, T. Brand, and B. Kollmeier, "The role of silent intervals for sentence intelligibility in fluctuating noise in hearing-impaired listeners," *International Journal of Audiology*, vol. 45, no. 1, pp. 26–33, jan 2006.
- [11] B. Kollmeier, A. Warzybok, S. Hochmuth, M. A. Zokoll, V. Uslar, T. Brand, and K. C. Wagener, "The multilingual matrix test: Principles, applications, and comparison across languages: A review," *International Journal of Audiology*, vol. 54, no. sup2, pp. 3–16, may 2015.
- [12] T. Brand and B. Kollmeier, "Efficient adaptive procedures for threshold and concurrent slope estimates for psychophysics and speech intelligibility tests," *The Journal of the Acoustical Society of America*, vol. 111, no. 6, pp. 2801–2810, jun 2002.
- [13] M. R. Schädler, A. Warzybok, S. Hochmuth, and B. Kollmeier, "Matrix sentence intelligibility prediction using an automatic speech recognition system," *International Journal of Audiology*, vol. 54, no. sup2, pp. 100–107, may 2015.
- [14] M. R. Schädler, D. Hülsmeier, G. Grimm, and , Maurice, "m-r-s/fade: Fade version 2.4.1," 2021. [Online]. Available: <https://zenodo.org/record/4730596>
- [15] D. Hülsmeier, M. Buhl, N. Wardenga, A. Warzybok, M. R. Schädler, and B. Kollmeier, "Inference of the distortion component of hearing impairment from speech recognition by predicting the effect of the attenuation component," *International Journal of Audiology*, pp. 1–15, jun 2021.