

VOICE QUALITY FEATURES FOR REPLAY ATTACK DETECTION

Abraham Woubie and Tom Bäckström

Department of Signal Processing and Acoustics, Aalto University, Espoo, Finland
abraham.zewoudie@aalto.fi, tom.backstrom@aalto.fi

ABSTRACT

Replay attacks are attempts to get fraudulent access to an automatic speaker verification system. In this paper, we investigate the usefulness of voice quality features to detect replay attacks. The voice quality features are used together with the state-of-the-art constant Q cepstral coefficients (CQCC) features. The two feature sets are fused at the score level. Thus, the log-likelihood scores estimated from the two feature sets are linearly weighted to obtain a single fused score. The fused score is used to classify whether a given speech sample is genuine or spoofed. Our experiments with the ASVspoof 2017 dataset demonstrate that the fusion of log-likelihood scores extracted from the CQCC and voice quality features improve the Equal Error Rate (EER) compared to the baseline system which is based only on CQCC features.

Index Terms— fusion, jitter, replay attack, shimmer, spoofing

1. INTRODUCTION

In recent years, the increasing availability of plenoptic cameras and research advances in light field capture, processing, and rendering, have been contributing for a growing interest in the extended capabilities offered by computational photography. Such capabilities include, for example, changing the point of view, refocusing, rendering all-in-focus images, and computing depth maps from the scene. Unlike traditional single-image acquisition systems, light fields enable promising advances in quite diverse application areas, such as multimedia for entertainment, medical, industry, science, etc

Automatic speaker verification is widely used in a range of applications which require not only robustness to changes in the acoustic environment, but also resilience to intentional circumvention, known as spoofing [21]. Spoofing is an attack where a fraudster tries to gain access of the system by masquerading as an enrolled person in the automatic speaker verification (ASV) system [16, 21]. The state-of-the-art ASV system are susceptible to different types of spoofing attacks such as voice conversion, speech synthesis, and replay attacks [5, 21].

Replay attacks are easy to perform and their threat to the reliability of ASV has been studied widely [1, 8, 22]. Replay

attacks use recordings of a target speaker's voice which is replayed to the ASV system in place of genuine speech [14, 19]. A prime example is to record and replay a target speaker's voice to unlock a smartphone which uses ASV for access control.

Recently, replay attacks attracted a lot of attention in the research community. For example, the ASVspoof 2017 challenge provided a standard corpora for combating replay attack [4]. The challenge uses the constant Q cepstral coefficients (CQCC) as the feature set and Gaussian mixture model (GMM) techniques as a classifier [4, 12]. In fact, most state-of-the-art of anti-spoofing systems, including [4], use CQCC as a feature set.

After the release of the the ASVspoof 2017 challenge, different types of features have been proposed by different researchers to improve the performance of anti-replay attacks. For instance, features such as spectral peak mapping filter cepstral coefficients, subband spectral centroid magnitude coefficients (SCMC), subband spectral centroid frequency coefficients (SCFC), Teager energy profiles and others have also been used to detect replay attacks [7, 9–11, 18].

Thus, the main contribution of this work is that we propose the use of voice quality features for anti-spoofing systems, particularly to detect replay attacks. The voice quality features are used together with the state-of-the-art constant Q cepstral coefficients (CQCCs) features. The voice quality features are fused with the CQCCs at the score likelihood level (i.e., the log-likelihood scores extracted using CQCC and voice quality models are linearly weighted). We are interested in voice quality features since jitter and shimmer measurements show significant differences between different speaking styles. In addition, since these features have shown potential for characterizing pathological voices and linguistic abnormalities, they can be also employed to characterize a particular speaker.

2. VOICE QUALITY FEATURES

voice quality features characterize the glottal excitation of signal of voiced voices such as glottal pulse shape and fundamental frequency, and carry speaker-specific information. Analysis of the voice quality of a person is a valuable technique for speech pathology detection [2, 23]. For example,




Fig. 1: Jitter measurements for 3 pitch periods




Fig. 2: Shimmer measurements for 3 pitch periods

voice-disorders can be analyzed using such acoustic signal parameters. Unlike the F_0 , voice quality features do not always have an acoustic characteristic that is easily distinguishable and measurable from a speech signal.

Jitter and shimmer voice quality features measure variations of the fundamental frequency and amplitude of pitch periods, respectively. They are very useful to describe the fluctuations of the voice signal in a qualitative way. They are given as a percentage that represents the maximum deviation from a normal frequency or amplitude. There are many possible jitter and shimmer measurements, but usually it is based on an auto-correlation method for determining the frequency and location of each cycle of vibration of the vocal folds (i.e., pitch marks) [17].

Jitter and shimmer voice quality features can be used to detect voice pathologies [20]. They are normally measured from long sustained vowels where voice quality measurement values outside a certain threshold are considered as pathological voices. In addition, voice quality features are related to the shape and dimension of the speaker’s vocal tract, and the way how the speech is generated by the voice production mechanism. Jitter and shimmer can also be used to characterize the age and the gender of a speaker [15]. Moreover, there are also significant differences in jitter and shimmer measurements between different speaking styles, especially in shimmer measurements [13]. Since these features have shown potential for characterizing pathological voices and linguistic abnormalities, they can be also employed to characterize a particular speaker.

There are many possible jitter and shimmer measurements. By using Praat [3], one can extract 5 different jitter and 6 different shimmer measurements.

Although there are different types of jitter and shimmer measurements as it is explained above, we have extracted 5

different jitter (and 4 different shimmer measurements encouraged by previous work of [6]. It is reported in [6] that these measurements provide better results for speaker recognition more than the other jitter and shimmer measurements. We have extracted the following types of measurements: Jitter (local), Jitter (local, absolute), Jitter (rap), Jitter (ppq5), Shimmer (local), Shimmer (local, dB), Shimmer (apq3) and Shimmer (apq11). A clear description of these measurements is found in [3].

3. PROPOSED SYSTEM

To improve the performance of the baseline anti-spoofing system, we propose a score-level framework that fuses the information provided by CQCC and voice quality features as it is shown in Fig. 3.

Firstly, the training data is partitioned into two sets: genuine and spoofed. Then, two types of features (i.e., CQCC and voice quality) are extracted from the genuine and spoofed data. Afterwards, we train two types of GMM models using genuine data: one GMM model using CQCC and another GMM model using voice quality features. Similarly, we train two types of GMM models using spoofed data: one GMM model using CQCC and another GMM model using voice quality features. Then, the log-likelihoods are predicted using the respective trained models for each feature set. Thus, four different log-likelihoods are computed using the genuine and spoofed models for CQCC and voice quality features. Finally, the log-likelihoods predicted for the voice quality are fused together with the log-likelihoods predicted using CQCC features. The GMM models are learned using expectation maximisation (EM) algorithm with random initialisation. Note that the baseline system uses only CQCC feature set.

Given an unseen test utterance, the CQCC and voice quality features are first computed. Then, they are scored with their respective models to obtain the log-likelihood scores. Afterwards, the two log-likelihood scores predicted using the two models are combined in a weighted fashion such that their weights sum to 1. Finally, the combined scores are used to make a decision (i.e., accept/reject a speech as genuine or spoofed).

The fused cosine-distance score is calculated as follows:

$$\begin{aligned} \Lambda(X) &= \alpha (\log L(X|\Theta_n) - \log L(X|\Theta_s)) + \\ &= ((1 - \alpha) (\log L(X|\vartheta_n) - \log L(X|\vartheta_s))), \end{aligned} \quad (1)$$

where $\Lambda(X)$ is the fused log-likelihood score, Θ_n and Θ_s are the genuine and spoofed GMM model using CQCC features, respectively and ϑ_n and ϑ_s are the genuine and spoofed GMM model using voice quality features, respectively. In addition, two different weights are applied on the predicted log-likelihood scores. While α weights the log-likelihood score predicted using CQCC, $(1 - \alpha)$ weights the log-likelihood scores from the voice quality features.

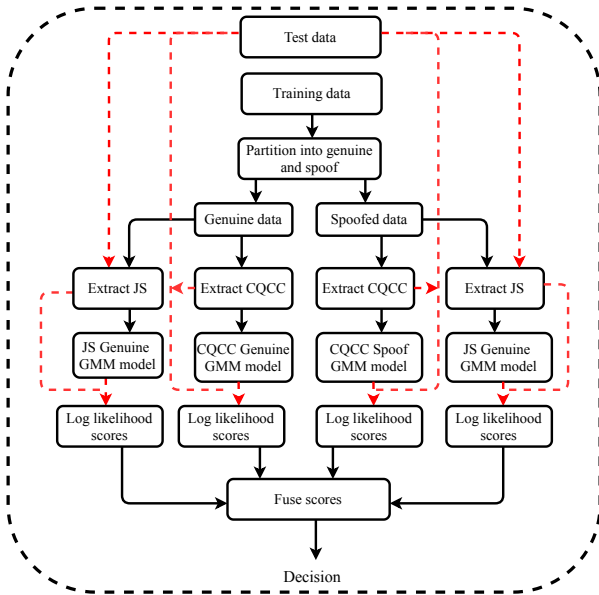


Fig. 3: The proposed replay attack detection system using CQCC and voice quality features. The baseline system is based only on CQCC features. While the arrows in black (undotted) correspond to training phase, the arrows in red (dotted) correspond to evaluation.

4. EXPERIMENTS

4.1. Experimental Setup

This work used the ASV Spoof 2017 challenge version 2.0 database [4]. The ASV Spoof 2017 database was collected in order to foster the development of countermeasures to protect ASV systems from replay spoofing attacks. It is partitioned into three subsets: training, development and evaluation. The number of files in the training, development and evaluation set are 3014, 1710 and 13306, respectively.

The baseline system [4] uses the CQCC feature set. The maximum and minimum frequency values are set to 8 kHz and 15 Hz, respectively. The number of bins per octave is 96 and the total number of CQCC features extracted include 19 static coefficients, log-energy, deltas, and delta-delta. Thus, the CQCC has a feature vector of length 60. After the extraction of CQCC features, the means and variances are normalized. The voice quality features are extracted over 30 ms frame length and at 10 ms shift using Praat [3]. Each of the voice quality features are then estimated over a 500 ms window with 10 ms shift. This is done to smooth out the feature estimation of the unvoiced frames. It is also done to synchronize the voice quality features with CQCC. Both the baseline and the proposed system uses GMM classifier for modeling the classes corresponding to natural and spoofed speech utterances. Since the size of CQCC feature is 60, we set its GMM

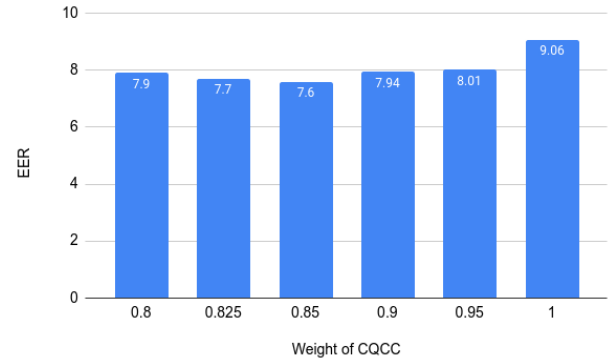


Fig. 4: Equal Error Rate (EER) of the development set when the weight of CQCC is tuned. The baseline system has weight value of 1.

components to 512. But, since the voice quality features has only 9 features, its GMM components is set to 128.

4.2. Experimental Results

We chose to use the work of [4] as a baseline system. We further selected the best results from [4] to compare the performance of the proposed system. The performance of the proposed and baseline systems were compared with three different experiments. The first system uses the training dataset (3710 files) as a training set and uses two test sets: development and evaluation. The second system uses the development set (1710 files) as training, and uses two test sets: training and evaluation. Finally, the training and development sets are pooled together (i.e., 4724 files) and are used as a training test, and the test is carried out on the evaluation set.

To find the best weight values for CQCC and voice quality features, we experimented with different values for the weight on the development set as it is shown in Fig. 4. The figure shows that the fusion of voice quality features with CQCC using different weight values provides better EER compared to the baseline system. Since the best EER value is found when the weight of CQCC is 0.85, we have used weight of 0.85 for CQCC and 0.15 for voice quality features for the results reported in Table 1.

From the results of Table 1, we see that the baseline system which uses the training data as a training set provides an EER of 9.06% and 13.74% on the development and evaluation sets, respectively. The table further shows that the fusion of voice quality features with CQCC gives an EER of 7.6% and 12.3% on the development and evaluation sets, respectively. Thus, the results show that the addition of voice quality measurements to the CQCC feature set provides a 16% and 10.48% relative EER improvement on the development and evaluation sets, respectively.

Similarly, Table 1 shows that when the development set is used as a training set, the baseline system provides an EER

Evaluation	Training				
	T		D		T + D
	D	E	T	E	E
CQCC [4]	9.06	13.74	5.66	14.77	12.24
CQCC + JS	7.6	12.3	4.01	13.2	10.9

Table 1: Replay detection performance in terms of Equal Error Rate (EER) for the ASVspoo 2017 Version 2.0 database for training (T), development (D) and testing (T) configurations. JS represents jitter and shimmer voice quality measurements. Note that the baseline system EER (i.e, CQCC) results are taken from the work of [4].

of 5.66% and 14.77% on the training and evaluation sets, respectively. The table further reveals that the augmentation of voice quality features with CQCC on the same dataset provides EER of 4.01% and 13.2% on the training and evaluation sets, respectively. These improvements represent a 29.15% and 10.62% relative EER improvement on the training and evaluation sets, respectively.

In a final experiment, we pooled the training and development files together and used them as a training set of 4724 files. The table shows that the baseline system provides an EER of 12.24%. However, the addition of voice quality features to the CQCC reduces the EER to 10.9%. This represents almost an 11% relative EER improvement compared to the baseline system.

The histogram plots of log-likelihood scores obtained from Gaussian mixtures corresponding to (a) CQCC, (b) voice quality features and (c) CQCC + voice quality features are shown in Fig. 5. The log-likelihood scores are for the evaluation set. From the figure, we see that the log-likelihood scores of the voice quality scores of both natural and replay are distributed more resulting in lower % EER as compared to the distribution obtained from CQCC. When the log-likelihood scores of CQCC and voice quality are fused, the scores of CQCC and voice quality are multiplied by 0.85 and 0.15, respectively.

In addition, the work in [4] also reported results of replay attack detection without mean variance normalization and without log energy. Thus, in order to further assess the impact of voice quality features, we have also made another set of experiments where the log-energy is not used at all and no normalization is carried out. Similar to the results reported in Table 1, our experimental results showed that the fusion of voice quality features with the CQCC, irrespective of log-energy and normalization, always provides better EER than the baseline system which is based only on CQCC features, respectively.

Thus, the results reported in Table 1 demonstrate that the long-term voice quality features provide useful and complementary speaker information. The experimental results show that adding jitter and shimmer voice quality features to the

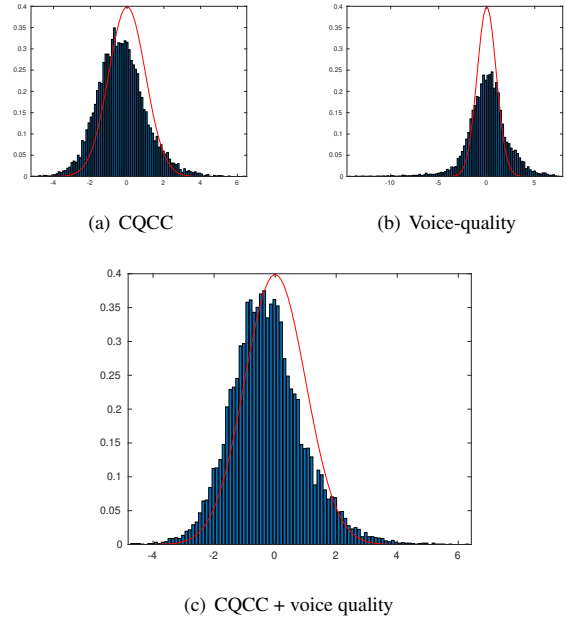


Fig. 5: Histogram of log-likelihoods scores of CQCC, voice quality, and CQCC + voice quality. The weights of CQCC and voice quality features are 0.85 and 0.15 in (c), respectively.

baseline CQCC features reduce the EER values. In overall, the use of voice quality features together with CQCC ones increase the robustness and reliability of anti-spoofing systems. Note that this work analyzed the robustness of voice quality to spoofing, not background noise.

5. CONCLUSIONS

In this work, we have proposed the use of jitter and shimmer voice quality measurements as a complementary source of information to detect replay attacks. The experimental results carried out on ASVspoo 2017 database demonstrate that the fusion of the log-likelihood scores of voice quality with the log-likelihood scores of CQCC improves the performance of anti replay attacks. The experimental results show that the augmentation of voice quality features with CQCC provide 11% relative EER improvement compared to using CQCC features on the evaluation set of ASVspoo 2017 database. Thus, the results reported in this work demonstrate the usefulness of voice quality measurements as a complementary source of information to detect replay attacks.

The future work could focus on applying deep neural network techniques on these long-term voice quality measurements to reduce the EER of anti-spoofing systems.

REFERENCES

- [1] F. Alegre, A. Janicki, and N. Evans. Re-assessing the threat of replay spoofing attacks against automatic speaker verification. In *2014 International Conference of the Biometrics Special Interest Group (BIOSIG)*, pages 1–6. IEEE, 2014.
- [2] S. Bielałowicz, J. Kreiman, B. R. Gerratt, M. S. Dauer, and G. S. Berke. Comparison of voice analysis systems for perturbation measurement. *Journal of Speech, Language, and Hearing Research*, 39(1):126–134, 1996.
- [3] P. Boersma and D. Weenink. Praat: doing phonetics by computer, 2009.
- [4] H. Delgado, M. Todisco, M. Sahidullah, N. Evans, T. Kinnunen, K. Lee, and J. Yamagishi. Asvspoof 2017 version 2.0: meta-data analysis and baseline enhancements. 2018.
- [5] N. Evans, J. Yamagishi, and T. Kinnunen. Spoofing and countermeasures for speaker verification: a need for standard corpora, protocols and metrics. *IEEE Signal Processing Society Speech and Language Technical Committee Newsletter*, pages 2013–05, 2013.
- [6] M. Farrús, J. Hernando, and P. Ejarque. Jitter and shimmer measurements for speaker recognition. In *Eighth annual conference of the international speech communication association*, 2007.
- [7] R. Font, J. M. Espín, and M. J. Cano. Experimental analysis of features for replay attack detection—results on the asvspoof 2017 challenge. In *Interspeech*, pages 7–11, 2017.
- [8] J. Gałka, M. Grzywacz, and R. Samborski. Playback attack detection for text-dependent speaker verification over telephone channels. *Speech Communication*, 67:143–153, 2015.
- [9] M. R. Kamble and H. A. Patil. Analysis of reverberation via teager energy features for replay spoof speech detection. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2607–2611. IEEE, 2019.
- [10] M. R. Kamble and H. A. Patil. Novel variable length teager energy profiles for replay spoof detection. *energy*, 32:33, 2020.
- [11] M. R. Kamble, H. Tak, and H. A. Patil. Effectiveness of speech demodulation-based features for replay detection. In *Interspeech*, pages 641–645, 2018.
- [12] T. Kinnunen, M. Sahidullah, H. Delgado, M. Todisco, N. Evans, J. Yamagishi, and K. A. Lee. The asvspoof 2017 challenge: Assessing the limits of replay spoofing attack detection. 2017.
- [13] X. Li, J. Tao, M. T. Johnson, J. Soltis, A. Savage, K. M. Leong, and J. D. Newman. Stress and emotion classification using jitter and shimmer features. In *Acoustics, Speech and Signal Processing, 2007. ICASSP 2007. IEEE International Conference on*, volume 4, pages IV–1081. IEEE, 2007.
- [14] J. Lindberg and M. Blomberg. Vulnerability in speaker verification—a study of technical impostor techniques. In *Sixth European Conference on Speech Communication and Technology*, 1999.
- [15] N. Minematsu, M. Sekiguchi, and K. Hirose. Automatic estimation of one’s age with his/her speech based upon acoustic modeling techniques of speakers. In *Acoustics, Speech, and Signal Processing (ICASSP), 2002 IEEE International Conference on*, volume 1, pages I–137. IEEE, 2002.
- [16] H. A. Patil and M. R. Kamble. A survey on replay attack detection for automatic speaker verification (asv) system. In *2018 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pages 1047–1053. IEEE, 2018.
- [17] J. Rusz, R. Cmejla, H. Ruzickova, and E. Ruzicka. Quantitative acoustic measurements for characterization of speech and voice disorders in early untreated parkinson’s disease. *The journal of the Acoustical Society of America*, 129(1):350–367, 2011.
- [18] W. Shang and M. Stevenson. Score normalization in playback attack detection. In *2010 IEEE international conference on acoustics, speech and signal processing*, pages 1678–1681. IEEE, 2010.
- [19] J. Villalba and E. Lleida. Speaker verification performance degradation against spoofing and tampering attacks. In *FALA workshop*, pages 131–134, 2010.
- [20] I. Wagner. A new jitter-algorithm to quantify hoarseness: an exploratory study. *International Journal of Speech Language and the Law*, 2(1):18–27, 2013.
- [21] Z. Wu, N. Evans, T. Kinnunen, J. Yamagishi, F. Alegre, and H. Li. Spoofing and countermeasures for speaker verification: A survey. *speech communication*, 66:130–153, 2015.
- [22] Z. Wu, S. Gao, E. S. Cling, and H. Li. A study on replay attack and anti-spoofing for text-dependent speaker verification. In *Signal and Information Processing Association Annual Summit and Conference (APSIPA), 2014 Asia-Pacific*, pages 1–5. IEEE, 2014.
- [23] I. C. Zwetsch, R. D. R. Fagundes, T. Russomano, and D. Scolari. Digital signal processing in the differential diagnosis of benign larynx diseases. *Scientia Medica*, 16(3):109–114, 2006.