

An explainability study of the constant Q cepstral coefficient spoofing countermeasure for automatic speaker verification

Hemlata Tak, Jose Patino, Andreas Nautsch, Nicholas Evans and Massimiliano Todisco

EURECOM, Sophia Antipolis, France

lastname@eurecom.fr

Abstract

Anti-spoofing for automatic speaker verification is now a well established area of research, with three competitive challenges having been held in the last 6 years. A great deal of research effort over this time has been invested into the development of front-end representations tailored to the spoofing detection task. One such approach known as constant Q cepstral coefficients (CQCCs) have been shown to be especially effective in detecting attacks implemented with a unit selection based speech synthesis algorithm. Despite their success, they largely fail in detecting other forms of spoofing attack where more traditional front-end representations give substantially better results. Similar differences were also observed in the most recent, 2019 edition of the ASVspoof challenge series. This paper reports our attempts to help explain these observations. The explanation is shown to lie in the level of attention paid by each front-end to different sub-band components of the spectrum. Thus far, surprisingly little has been learned about what artefacts are being detected by spoofing countermeasures. Our work hence aims to shed light upon signal or spectrum level artefacts that serve to distinguish different forms of spoofing attack from genuine, bone fide speech. With a better understanding of these artefacts we will be better positioned to design more reliable countermeasures.

Keywords: spoofing, presentation attack detection, constant Q cepstral coefficients, explainability

1. Introduction

Automatic speaker verification (ASV) technology is ubiquitous nowadays, being applied to user authentication in an increasingly diverse array of applications from smart-home technology to telephone banking to health-sector applications [1, 2]. Despite its success, there are justified concerns surrounding the vulnerabilities of ASV to spoofing, namely manipulation from inputs specially crafted to deceive the system and provoke false acceptances. The community has responded to this threat with spoofing countermeasures that aim to detect and deflect such attacks. The effort has been spearheaded through the ASVspoof initiative which emerged from the first special event on the topic held at INTERSPEECH in 2013 [3, 4].

The first edition of ASVspoof, held in 2015 [5], showed that some voice conversion and speech synthesis spoofing attacks can be detected with ease. The infamous S10 unit selection speech synthesis algorithm,¹ however, was shown to present substantial difficulties. The constant Q cepstral coefficient (CQCC) front-end [6, 7], introduced subsequently in 2016, was the first solution capable of detecting the S10 attack, reducing equal error rates (EERs) by some 72% relative to the state

of the art at the time. While better performing solutions have since emerged, the CQCC front-end remains popular. Over half of ASVspoof 2019 participants (26 out of 48) used CQCC representations as a component of their submission. This includes a number of top-performing systems, e.g. [8, 9, 10, 11].

CQCC features are, however, not a *silver bullet*. Even if near-to-zero EERs can be obtained for some attacks, for others results are poor. Among those in the ASVspoof 2015 database, for the S8 tensor-based approach to voice conversion [12], EERs obtained using CQCCs are substantially higher than those obtained with more conventional linear frequency cepstral coefficient (LFCC) features. These observations have led us to ask ourselves *why*? What could account for these performance differences and what can be learned from the explanation?

The answer lies in the artefacts that characterise spoofing attacks for which each front-end excels. Artefacts are signatures or tell tale signs in spoofed speech that are left behind by voice conversion and speech synthesis algorithms. They hence provide the means to distinguish between bona fide and spoofed speech. While end-to-end deep learning architectures, which operate directly upon the signal and which have the capacity to identify such artefacts automatically, may diminish the need to understand what makes one representation outperform another, explainability is *always* an issue of scientific interest.

While we have learned a great deal about the vulnerabilities of ASV to spoofing and the pattern recognition architectures that function well in detecting spoofing attacks, we have probably learned surprising little about what characterises spoofing attacks at the signal, spectrum or feature level. An understanding at this level would surely place us in a better position to design better performing countermeasures capable of detecting a broader range of attacks with greater reliability. The same understanding might also provide some assurance that the artefacts being detected are in fact indicative of spoofing attacks rather than database collection and pre-processing procedures.

The work presented in this paper hence revisits CQCC features and attempts to show why they are effective in detecting some attacks but less effective in detecting others. In particular, we are interested in shedding light upon what is being detected in the signal, i.e. the artefacts that characterise a spoofing attack. The paper is organised as follows. Section 2 presents a brief review of the CQCC front-end. Section 3 describes a sub-band approach to examine where information relevant to spoofing detection is located in the spectrum. Section 4 presents countermeasure performance and sub-band analysis results for the ASVspoof 2019 database. Section 5 describes the impact of variable spectro-temporal resolution whereas Section 6 presents additional experiments designed to validate our findings. A discussion of these is presented in Section 7, together with some directions for future work.

¹<http://mary.dfki.de/>

2. Constant Q cepstral coefficients

This section describes the motivation behind the development of constant Q cepstral coefficients (CQCCs). It shows how CQCC features are derived starting with the constant Q transform (CQT), describes differences to spectro-temporal decomposition with the discrete Fourier transform (DFT) and presents comparative results for the ASVspoof 2015 database.

2.1. Motivation

One strategy to spoof an ASV system would be to generate and present speech signals whose features match well those corresponding to genuine, bona fide speech produced by the target speaker (the identity being claimed by the attacker). In this sense, the characteristics of a speech signal that distinguish different speakers are unlikely to be the same as those that distinguish between bona fide and spoofed speech. These assumptions provided the motivation behind the development of CQCCs; instead of using features designed specifically for ASV, a popular choice in the early years of anti-spoofing research, we sought to use features that were fundamentally *different* to those designed for ASV. The fundamental assumption was that, if the spoofer is trying to produce speech signals that *replicate* the same features used for ASV, then features optimised for ASV cannot be the best approach to detect spoofed speech signals; features not optimised for ASV would have better potential for spoofing detection since the spoofer was not trying to produce speech signals which *replicate* these features. Although it was discovered later that CQCCs were also useful for other speech tasks including speaker diarization, ASV and utterance verification [7, 13], they were not *designed* for ASV. The inspiration for CQCCs came from prior experience of using the constant Q transform (CQT) for music processing tasks [14].

2.2. From the CQT to CQCCs

The perceptually motivated CQT [15, 16] approach to the spectro-temporal analysis of a discrete signal $x(n)$ is defined by:

$$X^{CQ}(k, n) = \sum_{l=n-\lfloor N_k/2 \rfloor}^{n+\lfloor N_k/2 \rfloor} x(j) a_k^*(l - n + N_k/2) \quad (1)$$

where n is the sample index, $k = 1, 2, \dots, K$ is the frequency bin index, $a_k(n)$ are the basis functions, $*$ is the complex conjugate and N_k is the frame length. The basis functions $a_k(n)$ are defined by:

$$a_k(n) = g_k(n) e^{j2\pi n f_k / f_s}, n \in \mathbb{Z} \quad (2)$$

where $g_k(n)$ is zero-centred window function, f_k is the center of each frequency bin and f_s is the sampling rate. Further details are available in [14].

The center of each frequency bin f_k is defined according to $f_k = 2^{(k-1)/(B)} f_1$, where f_1 is the center of the lowest frequency bin and B is the number of bins per octave. B determines the trade-off between spectral and temporal resolutions. $N_k \in \mathbb{R}$ in Eqs. 1 and 2 is real-valued and inversely proportional to f_k . As a result, the summation in Eq. 1 is over a number of samples that is dependent upon the frequency. The spectro-temporal resolution is hence also dependent upon the frequency. The quality (Q) factor, given by $Q = f_k / \delta f$, where δf is the bandwidth, reflects the selectivity of each filter in the filterbank. For the CQT transform, Q is constant for all frequency bins k ; filters are logarithmically spaced.

The CQT gives $X^{CQ}(k, n)$ a geometrically spaced spectrum. This is resampled using a spline interpolation method to a uniform, linear scale, giving $\bar{X}^{CQ}(l, n)$ which attributes equal weighting to information across the full spectrum [6]. CQCCs [6] are obtained from the discrete cosine transformation (DCT) of the logarithm of the squared-magnitude CQT as follows:

$$CQCC(p, n) = \sum_{l=0}^{L-1} \log \left| \bar{X}^{CQ}(l, n) \right|^2 \cos \left[\frac{p \left(l - \frac{1}{2} \right) \pi}{L} \right]$$

where $p = 0, 1, \dots, L - 1$ and where l is now the linear-scale frequency bin index. Full details of the CQCC extraction algorithm are available in [6]. Efficient implementations of the CQT can be found in [14] and [17].

2.3. Differences to the DFT and LFCC

The principal differences between the CQT and the discrete Fourier transform (DFT) relate to the spectro-temporal resolution. Whatever the algorithm, spectral decompositions essentially act as a filterbank. For the DFT, the set of filterbank frequencies are linearly distributed and the bandwidth of each filter is constant. For the DFT, Q is no longer constant, but is linearly distributed. The series of filters is no longer logarithmically spaced, but linearly spaced. This ensures that the DFT exhibits constant spectral resolution. This is not the case for the CQT. Compared to the DFT and as a direct consequence of the constant Q property, the CQT-derived spectrum has greater frequency resolution at lower frequencies than at higher frequencies. The CQT-derived spectrogram will then exhibit greater temporal resolution at higher frequencies than at lower frequencies.

2.4. Performance

Both CQCC and LFCC front-ends are typically used with simple Gaussian mixture model (GMM) classifiers. The performance obtained with GMM-CQCC and GMM-LFCC systems in terms of EER for the evaluation partition of the ASVspoof 2015 database is illustrated in Table 1, which reproduces results from [6]. Results are shown independently for each of the 10 different attacks. Results for two attacks show marked differences for the CQCC and LFCC front-ends. These are S8, for which LFCCs outperform CQCCs by 93% relative, and S10 for which CQCCs outperform LFCCs by 87% relative. While the motivation for our work stems from these two differences for the ASVspoof 2015 database, results reported later in this paper relate to the more recent ASVspoof 2019 database.

3. Sub-band analysis

The differences in the spectro-temporal resolution of the DFT and CQT lead to one potential explanation for differences in the performance of GMM-LFCC and GMM-CQCC systems. This relates to differences in the spectral resolution at lower frequencies and temporal resolution at higher frequencies. Put differently, it might suggest that the artefacts that distinguish spoofed speech from bona fide speech might reside in specific sub-bands, rather than in the full band signal. This hypothesis is supported by [18, 19, 20, 21] which all found this to be the case for synthetic speech. Accordingly, we set out to determine whether differences at the sub-band level could explain differences in the performance of LFCC and CQCC representations.

Table 1: EERs (%) for the ASVspoo 2015 database, evaluation partition. Results for GMM-CQCC and GMM-LFCC baseline systems, reproduced from [6]. Results for two attacks (S8 and S10) show stark differences in performance for each system.

System	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
GMM-CQCC	0.005	0.106	0.000	0.000	0.130	0.098	0.064	1.033	0.053	1.065
GMM-LFCC	0.027	0.408	0.000	0.000	0.114	0.149	0.011	0.074	0.027	8.185




Figure 1: A 2-dimensional heatmap visualisation of sub-band analysis results for an arbitrary spoofing attack. The horizontal axis depicts the high-pass cut-in frequency f_{\min} whereas the vertical axis depicts the low pass cut-off frequency f_{\max} . The colour bar to the right of the plot depicts the EER obtained for each band-pass filter configuration (pair of f_{\min} and f_{\max}).

Given the stark difference in temporal resolution at high frequencies, our hypothesis has always been that this is where the artefacts are located for attacks such as S10. The assumption is that these artefacts are then only captured by using a spectro-temporal decomposition with a higher temporal resolution.

3.1. Experiments

The set of experiments designed to test this hypothesis consist in an extensive sub-band analysis whereby the GMM-LFCC and GMM-CQCC classifiers are applied to the ASVspoo 2019 logical access database at a sub-band level. In any single experiment, the entire database is processed with a low-pass and/or high-pass filter. With both low-pass and high-pass filters, the result is a band-pass filter with cut-in f_{\min} and cut-off f_{\max} . Corresponding GMM models are retrained each time. Cut-in and cut-off frequencies are varied in steps of 400 Hz between 0 Hz and the sampling frequency of 8 kHz. The EER is then determined in the usual manner [22].

3.2. Heatmap visualisation

Results are visualised in the form of a 2-dimensional heatmap, an example of which is illustrated in Fig. 1. The horizontal vertical axes of the heatmap signify the cut-in frequency f_{\min} and cut-off frequency f_{\max} of the high-pass and low-pass filter respectively. For band-pass filters, $f_{\min} < f_{\max}$, hence the triangular form. The EER corresponding to each band-pass filter configuration is indicated with the colour bar to the right of Fig. 1, with blue colours indicating low EERs and red colours indicating higher EERs.

The left-most column of the heatmap shows the EER for

a decreasingly aggressive low-pass filter, moving from bottom-to-top. The top-most row of the heatmap shows the EER for an increasingly aggressive high-pass filter moving from left-to-right. Everywhere else corresponds to a band-pass filter, with the full-band configuration being at the top-left. EERs along the diagonal of Fig. 1 illustrate the benefit of spectrum information at the sub-band level.

The heatmap serves as a crude means to evaluate which sub-bands are the least and most informative in distinguishing between spoofed and bona fide speech. For the arbitrary example illustrated in Fig. 1, EERs along the diagonal suggest that information at lower frequencies is discriminative, whereas use of that at higher frequencies is not sufficiently discriminative to be used alone. EERs in the left-most column show that the most discriminative information of all is contained at low frequencies. So long as this information is used, then the EER is low. EERs in the top-most row show that as soon as low frequency information is discarded, then the EER increases. Information between 1 and 3 kHz and above 7 kHz is non-discriminative. That between 3 and 7 kHz is less discriminative, with reasonable EERs being obtained only when the information in a sufficient number of sub-bands is combined.

4. Experiments with the ASVspoo 2019 database

Described here is our use of the ASVspoo 2019 database, the specific configurations of GMM-LFCC and GMM-CQCC spoofing countermeasures, the subset of spoofing attacks with which this work was performed, and the results of sub-band analysis.

4.1. Database and protocols

The ASVspoo 2019 database consist of two different protocols based on logical access (LA) and physical access (PA) use case scenarios [23]. The LA scenario relates to spoofing attacks performed with voice conversion and speech synthesis algorithms whereas the PA scenario involves exclusively replay spoofing attacks. Since the work reported in this paper stems from observations made for spoofing attacks performed with voice conversion and synthetic speech algorithms, our experiments were performed with the LA dataset only.

All experiments were performed with the standard ASVspoo 2019 protocols which consist of independent train, development and, evaluation partitions. Spoofed speech in each partition is generated using a set of voice conversion and speech synthesis algorithms [24]. There are 19 different spoofing attacks. Attacks in the training and development set were created with a set of 6 algorithms (A01-A06), whereas those in the evaluation set were created with a set of 13 algorithms (A07-A19). All experiments reported in this paper were performed with the evaluation set only.

Table 2: EERs (%) for the ASVspoof 2019 logical access database, evaluation partition. Results for both baseline systems, GMM-CQCC (linear scale), GMM-LFCC and GMM-CQCC (geometric scale). Results for GMM-CQCC (linear) and GMM-LFCC reproduced from [23]. Results for six attacks (highlighted) show stark differences in performance for each system.

System	A07	A08	A09	A10	A11	A12	A13	A14	A15	A16	A17	A18	A19
GMM-CQCC (linear)	0.00	0.04	0.14	15.16	0.08	4.74	26.15	10.85	1.26	0.00	19.62	3.81	0.04
GMM-LFCC	12.86	0.37	0.00	18.97	0.12	4.92	9.57	1.22	2.22	6.31	7.71	3.58	13.94
GMM-CQCC (geometric)	3.39	0.34	0.46	6.86	4.62	3.58	4.23	0.67	1.52	4.00	25.04	19.63	29.46

Table 3: A summary of the six spoofing attack algorithms (voice conversion and speech synthesis) from the ASVspoof 2019 logical access database used in this work. * indicates neural networks. Full details for each algorithm can be found in [24].

Attack	Input	Input processor	Duration	Conversion	Speaker represent.	Outputs	Waveform generator
A07 [25],[26]	Text	NLP	RNN*	RNN*	One hot embed.	MCC, F0, BA	WORLD
A13 [27],[28]	Speech (TTS)	WORLD	DTW	Moment matching*	-	MCC	Waveform filtering
A14 [29],[30]	Speech (TTS)	ASR*	-	RNN*	-	MCC, F0, BAP	STRAIGHT
A16 [31]	Text	NLP	-	CART	-	MFCC, F0	Waveform concat.
A17 [32],[28]	Speech (human)	WORLD	-	VAE*	One hot embed.	MCC, F0	Waveform filtering
A19	Speech (human)	LPCC/MFCC	-	GMM-UBM	-	LPC	Spectral filtering + OLA

4.2. Countermeasures

Experiments were performed using the two spoofing countermeasures that were distributed with the ASVspoof 2019 challenge database. They are the GMM-LFCC and GMM-CQCC baseline systems² [23]. The front-end representation used by the GMM-LFCC system comprises 20 static, velocity (Δ) and acceleration ($\Delta\Delta$) LFCC coefficients, thereby giving 60-dimension feature vectors. The 90-dimension GMM-CQCC representation comprises 29 static with zero energy coefficient, Δ , $\Delta\Delta$ coefficients. The back-end is common to both systems and consists of a traditional GMM classifier with two models, one for bona fide speech and one for spoofed speech, both with 512 models. Scores are conventional log-likelihood ratios. Neither countermeasure system was re-optimised; this work is not concerned with optimisation. Results reported here are those obtained with the original ASVspoof 2019 baseline countermeasures as reported in [23].

4.3. Metric and performance

Even though it is not the primary metric for the ASVspoof 2019 database, performance is assessed in terms of the EER which is determined according to a convex hull approach [22]. While the community has moved towards use of the tandem detection cost function (t-DCF) metric [33], the primary metric for the ASVspoof 2019 database, the work reported in this paper is concerned exclusively with countermeasure performance rather than the impact of spoofing and countermeasures upon ASV performance. The latter is a related, but different issue. Use of the EER is then appropriate and sufficient for this purposes of the work presented in this paper.

EER results for the two countermeasures and the evaluation partition of the ASVspoof 2019 LA database are illustrated in Table 2. Just as is the case for the 2015 database, results for the 2019 database show substantial variations in performance for the GMM-LFCC and GMM-CQCC countermeasures. For attacks A07, A16 and A19, the GMM-CQCC system outperforms GMM-LFCC system whereas, for attacks A13, A14 and A17, it is the GMM-LFCC system that performs best, albeit still with relatively high error rates.

Subsequent experiments were performed separately for this subset of 6 specific spoofing attacks, all being examples of where one front-end representation leads to substantially better results than the other. Brief details of the specific algorithms used in creating each attack are illustrated in Table 3. There are four voice conversion algorithms (A13, A14, A17 and A19) and two speech synthesis algorithms (A07 and A16), even if the input to two of the voice conversion algorithms is also synthetic speech (A13 and A14). Further details of each algorithm are available in [24].

4.4. Sub-band analyses

We sought to explain differences in the performance of GMM-LFCC and GMM-CQCC systems using what we could learn from sub-band analysis, results for which are shown in Fig. 2. Heatmaps are shown for each of the 6 spoofing attacks: A07, A16, A19 for which the GMM-CQCC system performs best; A13, A14 and A17 for which the GMM-LFCC system performs best. Results for the GMM-CQCC and GMM-LFCC systems are illustrated in rows one and two respectively. In each case, the EER for the baseline system is denoted by the top-left-most point in each 2D heatmap, i.e. the full-band case.

Figs. 2 (a)-(c) show that the baseline GMM-CQCC system (top-left-most point) gives low EERs whereas Figs. 2 (g)-(i) show that the GMM-LFCC system gives consistently worse results. Figs. 2 (d)-(f) and (j)-(l) show the opposite, even if performance for A17 (l) is still poor for both the baseline systems. Upon inspection of heatmaps for both LFCC and CQCC front-ends, we see that the critical information for the detection of all three attacks lies at high frequencies. While it is not visible in the plots at this scale, the discriminative information lies above 7.6 kHz; near-to-zero EERs can be obtained using information between 7.6 and 8 kHz only, using either front-end.

The situation for attacks A13, A14 and A17 is a little different. The diagonals of the heatmaps in Figs. 2 (d), (e), (j) and (k) all suggest that the most informative information is at the lower frequencies. The left-most columns of Figs. 2 (d), (e) and (j) suggest that the most critical information is located at the very lowest sub-bands. For A17, neither front-end performs especially well. The CQCC front-end fails completely, whereas the LFCC front-end succeeds in capturing information across the majority of the full band.

²https://www.asvspoof.org/asvspoof2019/asvspoof2019_evaluation_plan.pdf

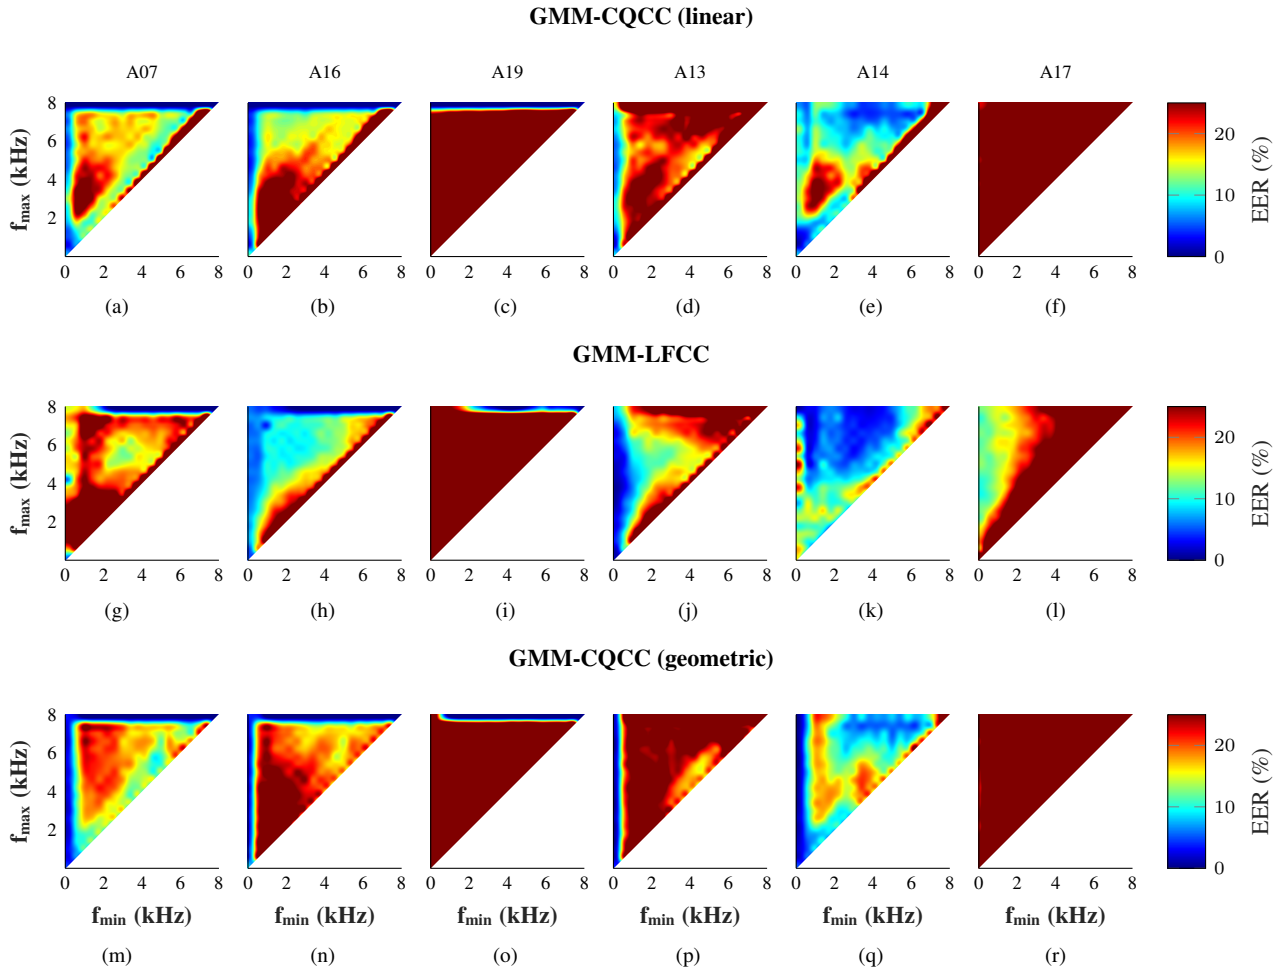


Figure 2: 2-D heatmap visualisations of sub-band analysis results for the six different spoof attacks. The first row (a-f) shows results for the GMM-CQCC (linear) system. The second row (g-l) shows corresponding results for the GMM-LFCC system. The third row (m-r) shows results for the GMM-CQCC (geometric) system.

5. Spectro-temporal resolution

Results presented above dispel somewhat the hypothesis that CQCCs produces better results than LFCCs for some attacks on account of the higher temporal resolution at high frequencies. If this were true, reliable performance would not have been obtained with the LFCC front-end using high frequency sub-bands alone. The explanation lies elsewhere.

The same results show that the discriminative information simply lies in the highest sub-bands; with appropriate band-pass filtering, both front-ends perform well for A07, A16 and A19 attacks. It is more than this, though, since such a straightforward explanation does not account for why the original, *full-band* CQCC front-end performs well, i.e. *without* band-pass filtering. The explanation for this observation requires us to revisit the issue of spectro-temporal resolution.

Fig. 3 illustrates a set of DFT and CQT-derived spectra for an arbitrary speech frame. Each plot also shows the second basis function of the DCT used in cepstral analysis. The vertical bars serve to illustrate the *spectral* resolution across the spectrum and give some indication of the *temporal* resolution, which is inversely proportional to the spectral resolution.

The top plot in Fig. 3 shows the CQT-derived spectrum (without re-sampling) [34]. It shows clearly that the spectral

resolution is higher at lower frequencies than at higher frequencies. The DCT basis function in this plot is *warped* to the same, non-linear frequency scale. The second plot shows exactly the same data, now plotted with a logarithmic frequency scale, hence the regularity of the vertical bars and DCT basis function. These plots show how, without resampling, the DCT will attribute greater *weight* to the low frequency components than to high frequency components. The third plot shows the CQT-derived spectrum after resampling. The vertical bars show how resampling acts to linearise the rate of sampling in the frequency domain. Note the difference in frequency scales for the second and third plots. The higher spectral resolution (and hence lower temporal resolution) for lower frequencies is still clearly apparent; the spectrum is much smoother at higher frequencies.

A comparison of the cosine basis functions to the CQT-derived spectra in the second and third plots shows that re-sampling acts to dilute (weaken) information at lower frequencies but to distil (emphasise) information at higher frequencies. Hence, spoofing artefacts at high frequencies are emphasised by the CQCC front-end.

The fourth plot of Fig. 3 shows the DFT-derived spectrum and the linear sampling in the frequency domain. Here, the DCT

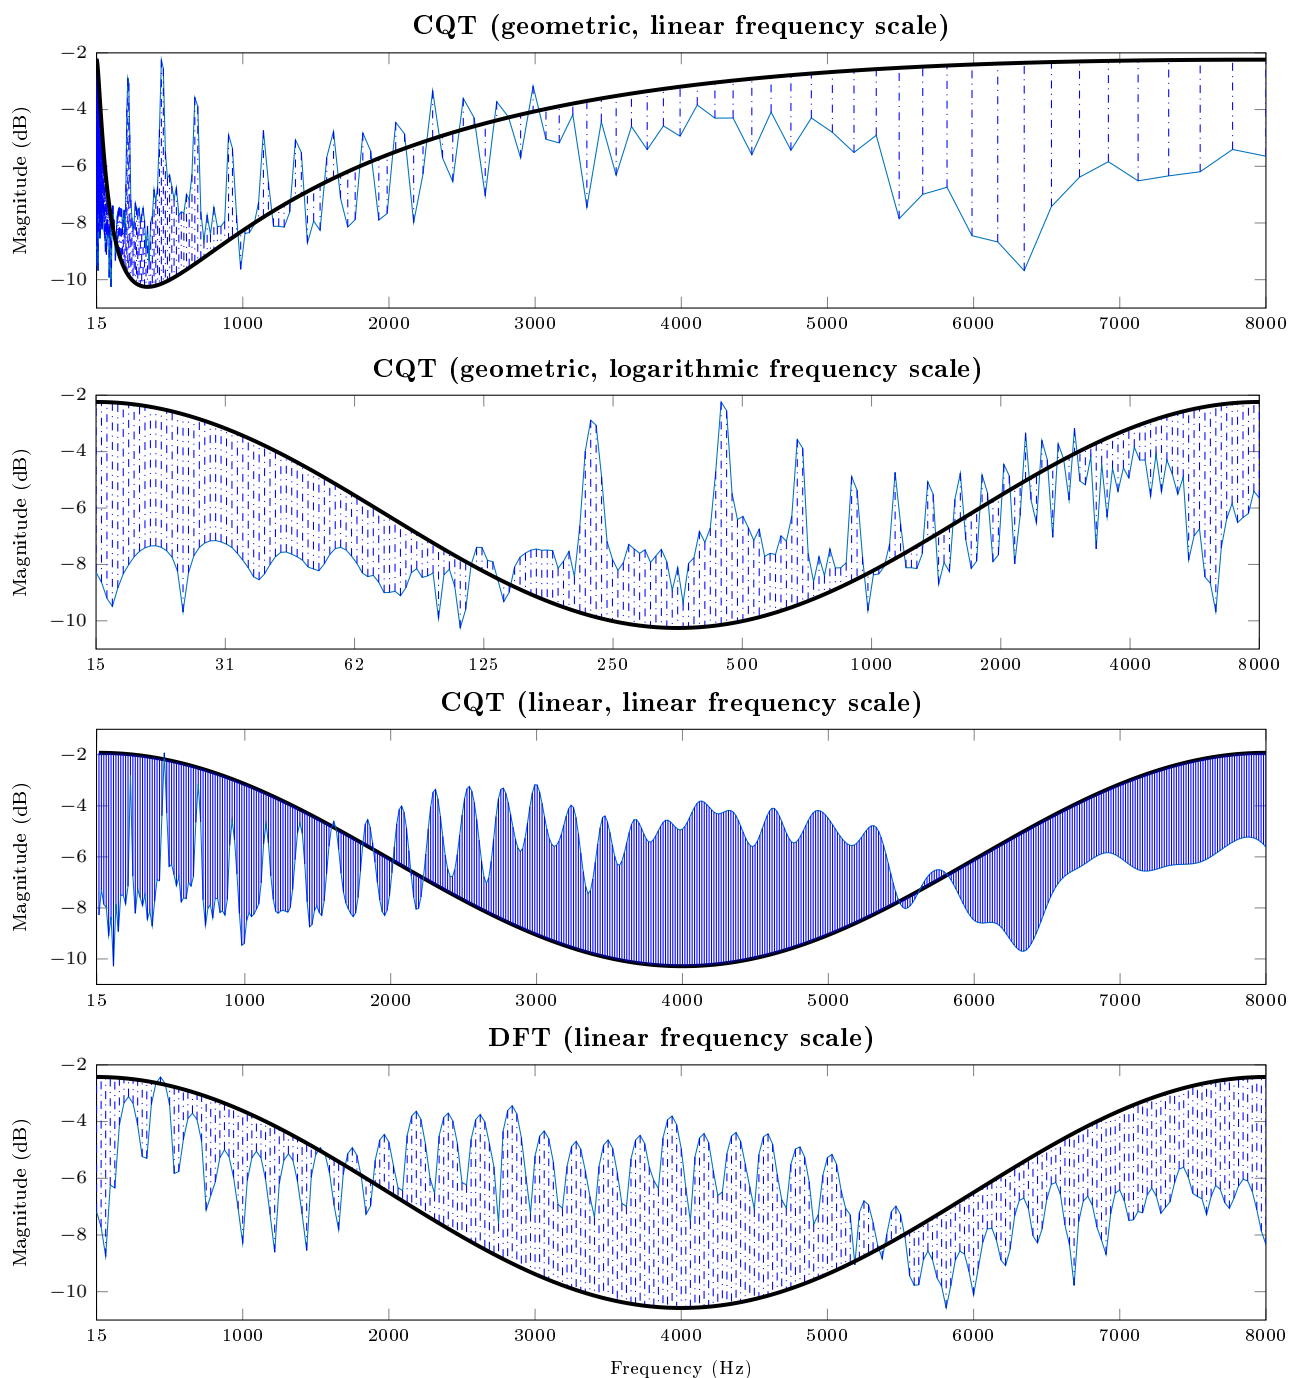


Figure 3: Illustrations of spectra for an arbitrary speech frame derived using the CQT and DFT. The top plot shows the CQT-derived spectrum without resampling, hence the logarithmic sampling in the linear frequency domain. The second plot shows the same spectrum but with a warped, logarithmic frequency scale. The third plot shows the CQT-derived spectrum after resampling to a uniform, linear frequency scale (Section 2). The fourth plot shows the corresponding DFT-derived spectrum. Also illustrated in each plot is the second basis function of the discrete cosine transform used in cepstral analysis.

acts to weight all frequency components uniformly. Spoofing artefacts at high frequencies are not emphasised; reliable performance is then obtained only by band-pass filtering.

The explanation for why the LFCC front-end outperforms the CQCC front-end for attacks A13, A14 and A17 is now straightforward. The artefacts for these attacks are at lower and

mid-range frequencies where the CQT acts to dilute information, hence why the GMM-CQCC system performs poorly for these attacks. Since the DFT gives uniform weighting to information at all frequencies, the GMM-LFCC system gives better, though still somewhat poor performance. This is because information at low frequencies is not *emphasised*.

6. Validation

We designed one final experiment to validate our findings. The CQT-derived spectra in the top two plots of Fig. 3 show dense sampling of the spectrum at lower frequencies but sparse sampling at higher frequencies. Hence, without resampling, the application of cepstral analysis should result in the emphasis of information situated at lower frequencies. In this case, a CQCC front-end *without* linear resampling should produce better detection performance in the case of attacks A13 and A14. Performance for attack A17 might not be improved, since the artefacts appear not to be localised at low frequencies.

We repeated the sub-band analysis described in Section 3 using the geometrically-scaled CQCC front-end. Results are illustrated in Figs. 2 (m)-(r) for the same set of 6 attacks for which results are described in Section 4. The comparison of the right-most columns in the heatmaps of Figs. 2 (d) and (p) and those in (e) and (q) clearly show that lower EERs are obtained using the geometrically-scaled rather than the linearly-scaled CQCC front-end. Low EERs are even obtained when the front-end is applied without any band-pass filtering. This is because use of the original CQT geometric frequency scale results in the emphasis of information at lower frequencies where the artefacts are located. As expected, attack A17 remains troublesome.

EERs for the geometrically-scaled CQCC front-end are presented in the last row of Table 2. For some attacks, the EER for the geometrically-scaled CQCC front-end is slightly higher. For some others, it is slightly lower. Once again, performance for A17 is poor, and even worse than for both the linearly-scaled CQCC and LFCC front-ends. Those for A13 and A14 clearly confirm our findings, with substantially lower EERs than those obtained with both the linearly-scaled CQCC and LFCC front-ends. Use of the original geometric scale of the CQT clearly results in the emphasis of information at lower frequencies where the artefacts are localised.

7. Discussion

This paper reports an explainability study of constant Q cepstral coefficients (CQCCs), a spoofing countermeasure front-end for automatic speaker verification [6, 7]. The work aims to explain why the CQCC front-end works so reliably in detecting some forms of spoofing attack, but performs poorly in detecting others. It confirms that different spoofing attacks exhibit artefacts at different frequencies, artefacts that are better captured with specific front-ends. The standard constant Q transform (CQT) [15, 16] exhibits a dense sampling of the spectrum at lower frequencies and a sparse sampling at higher frequencies. Hence, geometrically-sampled CQCCs perform well in detecting spoofing artefacts when they are located at low frequencies. Linear sampling shifts the emphasis to higher frequencies so that the detection of spoofing artefacts at high frequencies are emphasised and captured reliably.

Taken together, the results presented in this paper show that no single CQCC front-end configuration can perform well for all spoofing attacks; different spoofing attacks produce artefacts at different parts of the spectrum and these can only be detected reliably when the front-end emphasises information in the relevant frequency bands. This finding certainly explains why classifier fusion has proven to be so important to generalisation, i.e. reliable performance in the face of varying spoofing attacks. Given that spoofing artefacts can be highly localised in the spectrum, the same finding suggests that we may need to rethink the use of cepstral analysis, which acts to smooth information lo-

cated across the full spectrum. Approaches to detect localised artefacts need further investigation. We may then need to rethink the approach to classifier fusion too.

One natural extension of the work would be to look deeper into the artefacts themselves. Preliminary investigations show that those at high frequencies appear to be nothing more than band-limiting effects introduced by the spoofing algorithms, or the data used in their training. The source of those at lower frequencies are less clear. A clearer understanding of the artefacts at the signal level, rather than just the feature or spectrum level is certainly needed. This could help us to design better spoofing countermeasures, even if the same understanding might also help the fraudsters to design better spoofing attacks.

Lastly, while it was not the objective of this work to investigate the generalisation capabilities of the CQCC front-end, it is clear that none of the countermeasures for which results are reported in this paper is successful in detecting *all* of the spoofing attacks used in the creation of the ASVspoof 2019 database. Attacks A10, A12, A13, A17 and A18 are all examples where *neither* configuration of the CQCC front-end, *nor* the LFCC front-end are especially effective. The work presented in this paper focused on countermeasure performance exclusively. It did not consider the effect of spoofing and countermeasures upon automatic speaker verification (ASV) performance. Some attacks, e.g. A17, are known not to be especially effective in fooling ASV, and in this sense they pose no threat; that countermeasures fail to detect such attacks is hence of no serious concern. Other attacks, though, are effective in manipulating ASV and are also difficult to detect reliably. Other work reported by the community has shown far better results than those reported in this paper [8, 9, 10, 11]. Still, even the very best approaches can be vulnerable to certain, specific attacks. The *silver bullet* remains elusive for the time being.

8. References

- [1] J. Hansen and T. Hasan, "Speaker recognition by machines and humans: A tutorial review," in *IEEE Signal Processing Magazine*, vol. 32, no. 6, pp. 74–99, 2015.
- [2] T. Kinnunen and H. Li, "An overview of text-independent speaker recognition: From features to supervectors," *Speech communication*, vol. 52, no. 1, pp. 12–40, 2010.
- [3] N. Evans, T. Kinnunen, and J. Yamagishi, "Spoofing and countermeasures for automatic speaker verification," in *Proc. INTERSPEECH*, Lyon, FRANCE, 2013, pp. 925–929.
- [4] 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*, 2013.
- [5] Z. Wu, T. Kinnunen, N. Evans, J. Yamagishi, C. Hanilci, M. Sahidullah, and A. Sizov, "ASVspoof 2015: the first automatic speaker verification spoofing and countermeasures challenge," in *Proc. INTERSPEECH*, Dresden, Germany, 2015, pp. 2037–2041.
- [6] M. Todisco, H. Delgado, and N. Evans, "A new feature for automatic speaker verification anti-spoofing: Constant Q cepstral coefficients," in *Proc. Speaker Odyssey Workshop*, Bilbao, Spain, 2016, pp. 249–252.
- [7] M. Todisco, H. Delgado, and N. Evans, "Constant Q cepstral coefficients: A spoofing countermeasure for au-

- omatic speaker verification,” *Computer Speech & Language*, vol. 45, pp. 516–535, 2017.
- [8] B. Chettri, D. Stoller, V. Morfi, M. A. Martínez Ramírez, E. Benetos, and B. L. Sturm, “Ensemble models for spoofing detection in automatic speaker verification,” in *Proc. INTERSPEECH*, Graz, Austria, 2019, pp. 1118–1112.
- [9] R. Alluri and A. Vuppala, “IIIT-H spoofing countermeasures for automatic speaker verification spoofing and countermeasures challenge 2019,” in *Proc. INTERSPEECH*, Graz, Austria, 2019, pp. 1043–1047.
- [10] Y. Yang, H. Wang, H. Dinkel, Z. Chen, S. Wang, Y. Qian, and K. Yu, “The SJTU robust anti-spoofing system for the ASVspoof 2019 challenge,” in *Proc. INTERSPEECH*, Graz, Austria, 2019, pp. 1038–1042.
- [11] G. Lavrentyeva, S. Novoselov, A. Tseren, M. Volkova, A. Gorlanov, and A. Kozlov, “STC antispoofing systems for the ASVspoof2019 challenge,” in *Proc. INTERSPEECH*, Graz, Austria, 2019, pp. 1033–1037.
- [12] D. Saito, K. Yamamoto, N. Minematsu, and K. Hirose, “One-to-many voice conversion based on tensor representation of speaker space,” in *Proc. INTERSPEECH*, Florence, Italy, 2011, pp. 653–656.
- [13] T. Kinnunen, M. Sahidullah, I. Kukanov, H. Delgado, M. Todisco, A. Sarkar, N. B. Thomsen, V. Hautamäki, N. Evans, and Z. Hua Tan, “Utterance verification for text-dependent speaker recognition: a comparative assessment using the RedDots corpus,” in *Proc. INTERSPEECH*, San Francisco, USA, 2016, pp. 430–434.
- [14] C. Schörkhuber, A. Klapuri, N. Holighaus, and M. Dörfler, “A matlab toolbox for efficient perfect reconstruction time-frequency transforms with log-frequency resolution,” in *Proc. Audio Engineering Society International Conference on Semantic Audio*, London, UK, 2014.
- [15] J. Youngberg and S. Boll, “Constant-q signal analysis and synthesis,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Tulsa, Oklahoma, USA, 1978, vol. 3, pp. 375–378.
- [16] J. Brown, “Calculation of a constant Q spectral transform,” *Journal of the Acoustical Society of America (JASA)*, vol. 89, no. 1, pp. 425–434, 1991.
- [17] G. A. Velasco, N. Holighaus, M. Dörfler, and T. Frill, “Constructing an invertible constant-Q transform with nonstationary Gabor frames,” in *Proc. Digital Audio Effects (DAFx-11)*, Paris, France, 2011, pp. 93–99.
- [18] T. F. Quatieri, *Discrete-time speech signal processing: principles and practice*, Pearson Education, 2006.
- [19] J. R. Deller Jr, J. G. Proakis, and J. H. Hansen, *Discrete time processing of speech signals*, Prentice Hall PTR, 1993.
- [20] K. Sriskandaraja, V. Sethu, P. Ngoc Le, and E. Ambikairajah, “Investigation of sub-band discriminative information between spoofed and genuine speech,” in *Proc. INTERSPEECH*, San Francisco, USA, 2016, pp. 1710–1714.
- [21] J. Yang, R. K. Das, and H. Li, “Significance of subband features for synthetic speech detection,” *IEEE Transactions on Information Forensics and Security*, 2019.
- [22] N. Brümmer and E. De Villiers, “The bosaris toolkit: Theory, algorithms and code for surviving the new dcf,” *arXiv preprint arXiv:1304.2865*, 2013.
- [23] M. Todisco, X. Wang, V. Vestman, M. Sahidullah, H. Delgado, A. Nautsch, J. Yamagishi, N. Evans, T. Kinnunen, and K. Lee, “ASVspoof 2019: Future horizons in spoofed and fake audio detection,” in *Proc. INTERSPEECH*, Graz, Austria, 2019, pp. 1008–1012.
- [24] X. Wang, J. Yamagishi, M. Todisco, H. Delgado, A. Nautsch, N. Evans, M. Sahidullah, V. Vestman, T. Kinnunen, K. Lee, et al., “The ASVspoof 2019 database,” *arXiv preprint arXiv:1911.01601*, 2019.
- [25] K. Tanaka, H. Kameoka, and N. Kaneko, T. and Hojo, “Wavecyclegan2: Time-domain neural post-filter for speech waveform generation,” *arXiv preprint arXiv:1904.02892*, 2019.
- [26] Z. Wu, O. Watts, and S. King, “Merlin: An open source neural network speech synthesis system.,” in *Speech synthesis workshop (SSW)*, Sunnyvale, USA, 2016, pp. 202–207.
- [27] Y. Li, K. Swersky, and R. Zemel, “Generative moment matching networks,” in *International Conference on Machine Learning*, Lille, France, 2015, pp. 1718–1727.
- [28] K. Kobayashi, T. Toda, and S. Nakamura, “Intra-gender statistical singing voice conversion with direct waveform modification using log-spectral differential,” *Speech Communication*, vol. 99, pp. 211–220, 2018.
- [29] L. Juan Liu, Z. Hua Ling, Y. Jiang, M. Zhou, and L. Rong Dai, “WaveNet vocoder with limited training data for voice conversion,” in *Proc. INTERSPEECH*, Hyderabad, India, 2018, pp. 1983–1987.
- [30] H. Kawahara, I. Masuda-Katsuse, and A. De Cheveigné, “Restructuring speech representations using a pitch-adaptive time-frequency smoothing and an instantaneous-frequency based F0 extraction: Possible role of a repetitive structure in sounds,” *Speech Communication*, vol. 27, no. 3–4, pp. 187–207, 1999.
- [31] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, “Montreal Forced Aligner: Trainable text-speech alignment using Kaldi,” in *Proc. INTERSPEECH*, Stockholm, Sweden, 2017, pp. 498–502.
- [32] W. Chin Huang, Y. Chiao Wu, K. Kobayashi, Y. Huai Peng, H. Te Hwang, P. Lumban Tobing, Y. Tsao, H. Min Wang, and T. Toda, “Generalization of spectrum differential based direct waveform modification for voice conversion,” in *Proc. Speech synthesis workshop (SSW)*, Vienna, Austria, 2019, pp. 57–62.
- [33] T. Kinnunen, K. Lee, H. Delgado, N. Evans, M. Todisco, J. Sahidullah, M. and Yamagishi, and D. A. Reynolds, “t-DCF: a detection cost function for the tandem assessment of spoofing countermeasures and automatic speaker verification,” in *Proc. Speaker Odyssey Workshop*, Les Sables d’Olonne, France, 2018, pp. 312–319.
- [34] J. Brown, “Computer identification of musical instruments using pattern recognition with cepstral coefficients as features,” *The Journal of the Acoustical Society of America (JASA)*, vol. 105, no. 3, pp. 1933–1941, 1999.