

Maximum Voiced Frequency Estimation: Exploiting Amplitude and Phase Spectra

Thomas Drugman, *Member, IEEE*, Yannis Stylianou, *Senior Member, IEEE*

Abstract—Maximum Voiced Frequency (MVF) is used in various speech models as the spectral boundary separating periodic and aperiodic components during the production of voiced sounds. Recent studies have shown that its proper estimation and modeling enhance the quality of statistical parametric speech synthesizers. Contrastingly, these same methods of MVF estimation have been reported to degrade the performance of singing voice synthesizers. This paper proposes a new approach for MVF estimation which exploits both amplitude and phase spectra. It is shown that phase conveys relevant information about the harmonicity of the voice signal, and that it can be jointly used with features derived from the amplitude spectrum. This information is further integrated into a maximum likelihood criterion which provides a decision about the MVF estimate. The proposed technique is compared to two state-of-the-art methods, and shows a superior performance in both objective and subjective evaluations. Perceptual tests indicate a drastic improvement in high-pitched voices.

Index Terms—Maximum Voiced Frequency, Phase Processing, Speech Analysis, Speech Synthesis, Singing Voice, High-pitched.

I. INTRODUCTION

During the production of voiced sounds, a variety of models assume the speech signal to contain both periodic and aperiodic components. Two main strategies have been proposed in the literature to control the spectral weighting of these two components. The first one relies on a multiband approach where, for each spectral band, the energy ratio between the periodic and aperiodic contributions is controlled by *aperiodicity measurements*. These measurements can be computed in various ways. In [1], they consist of correlation coefficients calculated in each band, while in [2], [3] they are determined based on the ratio between the upper and lower smoothed spectral envelopes. The second approach for spectral weighting assumes the spectrum to be split into two bands: periodic components hold only in the low frequencies, while the aperiodic contributions take place in the higher part of the spectrum. The boundary between these two spectral bands is called the *maximum voiced frequency* (MVF), and it has been used in various speech models such as the so-called multiband excitation vocoder [4], the Harmonic plus Noise Model (HNM, [5], [6]) and its variants [7], the Deterministic plus Stochastic Model (DSM, [8]) of the excitation signal.

In the early versions of these vocoders, a fixed constant value of the MVF (generally around 4 kHz) was used for a particular speaker with a given voice quality [9], [5], [10].

Recent studies however have reported a gain by considering an explicit dynamic modeling of the MVF in Hidden Markov Model (HMM) based speech synthesis [11], [6]. In [11], the use of a dynamic MVF was even found to be slightly preferred over the multiband approach. While current methods of MVF estimation seem to be efficient in speech, some issues were reported in [12] for the synthesis of singing voices. For high-pitched voices, MVF was observed to be underestimated which led to an excessive amount of noise after synthesis.

All of these studies motivate the need for a technique which estimates accurately the MVF contour both in speech and singing voices. Very few methods have been proposed for this purpose. In [5], a peak-to-valley (P2V) measure is calculated for all possible harmonic candidates. In [6], the authors make use of the Sinusoidal Likeness Measure (SLM), which can be seen as a localized cross-correlation between the harmonic peak and the spectrum of a pure sinusoid. Based on these measurements for each harmonic candidate, the MVF decision can be taken using various strategies. In [5], a binary decision (into harmonic/non-harmonic) is taken for each spectral peak by applying a threshold to the P2V values, and the MVF is defined as the highest frequency for which the harmonicity criterion is met. In [6], SLMs are scaled between 0 and 1 by a nonlinear function and the MVF is defined by minimizing the Euclidian distance to the ideal SLM profile (i.e. 1s below the MVF and 0s beyond it). Since the resulting contours generally exhibit rapid irrelevant variations, a further post-processing is applied: median filtering in [5] and dynamic programming search in [6]. In [13], the authors proposed to replace the block of MVF determination based on the binary decisions in the P2V framework, by proposing various other strategies. They however reported a degradation of the perceptual quality after synthesis. Finally, it was proposed in [14] to classify spectral peaks as sine or noise based on 3 features describing their bandwidth, duration and frequency coherence.

The goal of this paper is to propose new techniques of MVF estimation. With regard to the state of the art, the contributions of this paper are the following: *i*) we investigate if measurements made directly on phase spectra contain relevant information for MVF estimation, *ii*) we envisage various joint uses of amplitude and phase-based measurements, *iii*) we propose a Maximum Likelihood (ML) criterion as a strategy to derive the MVF decisions, *iv*) contrary to current techniques which make use of some *arbitrary* settings, the proposed method is based on empirical settings derived from an objective study carried out on semi-synthetic signals.

The paper is organized as follows. The proposed approach is detailed in Section II. Both objective (Section III) and

subjective (Section IV) evaluations compare the performance of the proposed technique with those of P2V and SLM methods. Section V finally concludes the paper.

II. THE PROPOSED TECHNIQUE

The workflow of the proposed technique is presented in Figure 1. The steps involved in this workflow are further detailed in the following sections.




Fig. 1. Workflow of the proposed method.

A. Windowing

Windowing is essential as it determines the harmonicity properties of the resulting spectra. In all cases, the window length should be proportionnal to the pitch period. In this work, we have used a 4 period-long Hanning window as we found it to be suited for the amplitude spectra to exhibit a good peak-to-valley structure.

B. Spectral Peak Detection

Starting from an initial estimate of the fundamental frequency F_0 , the goal of this block is to find the angular frequencies ω_p of the harmonic candidates up to the Nyquist frequency. A standard peak picking is here applied, where ω_p is defined as the maximum value in the amplitude spectrum in the neighborhood (10 Hz on each side) of $p \cdot \omega_0$. Then ω_0 is updated to $\frac{\omega_p}{p}$ for the following harmonics.

C. Feature Extraction

For each harmonic candidate ω_p , we propose to extract three types of measurements x_p : Amplitude Spectrum (AS) based, Inter-Harmonic Phase Coherence (IHPC) and Inter-Cycle Phase Coherence (ICPC). The AS feature is close to the P2V measure proposed in [5]. For a harmonic candidate in ω_p , we consider the AS in the $[\omega_p - \frac{\omega_0}{2}; \omega_p + \frac{\omega_0}{2}]$ span. A local Harmonic-to-Noise Ratio (HNR) is then estimated by calculating the subtraction (in dB) between the average level in $[\omega_p - \frac{\omega_0}{5}; \omega_p + \frac{\omega_0}{5}]$, as it corresponds approximately to the main lobe, and the average level in the rest of the initial span.

The idea of IHPC is that the phase values across two consecutive harmonics should be similar, while this should not be the case for noisy contributions in speech. For this purpose, we use the Group Delay (GD) function, defined as $\frac{-d\phi(\omega)}{d\omega}$ where $\phi(\omega)$ is the unwrapped DFT phase spectrum. Note that a sufficiently high number of DFT points has to be used to facilitate phase unwrapping. The proposed IHPC is then expressed as the difference $GD(\omega_{p+1}) - GD(\omega_p)$ between two consecutive harmonic candidates.

The idea behind ICPC is to exploit the phase coherence for a particular harmonic when the windowing is shifted in time by the pitch period (T_0). Let us denote $\phi_1(\omega)$ the phase spectrum in the current frame, and $\phi_2(\omega)$ the phase spectrum

after applying a delay of T_0 (and compensating the linear phase due to this delay). The differential phase is defined as $\Delta\phi(\omega) = \phi_1(\omega) - \phi_2(\omega)$. The ICPC feature which is used throughout this paper is expressed as the wrapped value of $\Delta\phi(\omega_{p+1}) - \Delta\phi(\omega_p)$.

Other formulations of the AS, IHPC and ICPC features are obviously possible. According to our experiments, however, the expressions proposed in this work led to a robust and discriminative characterization of the harmonicity at a given spectral peak. As an illustration, Figure 2 shows the distributions of these 3 features for harmonic candidates extracted on a development set of semi-synthetic signals (see Section III-A) for which the actual decision of harmonicity is known. It can be seen that the proposed features convey a relevant information to predict such a decision. It is also worth noting that the 3 proposed features are invariant to energy scaling.




Fig. 2. Feature distributions for the 3 proposed features: AS (top plot), IHPC (middle plot) and ICPC (bottom plot). Histograms for real harmonics are in green bold line, while those of non-harmonics are in red dashed line.

D. Decision Strategy

To find the most likely MVF, we consider the null hypothesis H_0 as a rejection of the harmonicity criterion. The estimated MVF is then:

$$MVF = \underset{\omega_m}{\operatorname{argmax}} \left[\prod_{k=1}^m p(H_1|x_k, \omega_k) \cdot \prod_{l=m+1}^{N_{\text{harm}}} p(H_0|x_l, \omega_l) \right] \quad (1)$$

where x_p and N_{harm} respectively denote the measurement value and the number of harmonic candidates over the full band. The posterior for hypothesis H_i can be written as:

$$p(H_i|x_p, \omega_p) = \frac{p(x_p|H_i, \omega_p) \cdot p(H_i|\omega_p)}{\sum_{j=0}^1 p(x_p|H_j, \omega_p) \cdot p(H_j|\omega_p)} \quad (2)$$

Since the dependencies in ω_p are highly variable with the voice quality, speaker identity, recording conditions or voice type (e.g. speech vs singing voice), they will be omitted in the following. Considering both hypotheses to be equiprobable, it can be easily shown that (1) merely translates to the following ML criterion:

$$MVF = \underset{\omega_m}{\operatorname{argmax}} \left[\prod_{k=1}^m p(x_k|H_1) \cdot \prod_{l=m+1}^{N_{\text{harm}}} p(x_l|H_0) \right] \quad (3)$$

In the rest of this paper, the strategy decision was made according to (3) where likelihoods were calculated by a Gaussian approximation of the feature distributions shown in Figure 2 (and which were obtained on the development set described in Section III-A). Note that in practice we worked with log-likelihoods for computational precision reasons.

E. Time Smoothing

As for the P2V and SLM measures, the obtained MVF trajectories might contain spurious values and require some post-processing. For the objective evaluation (Section III), we applied a median filter of order 5 (considering a frame shift of 10 ms) in order to have results comparable with P2V. For the subjective evaluation, we preferred to use a moving average filter with a time constant of 30 ms (on both sides), since median filtering could lead to important jumps in the estimated contours, which could possibly lead to some auditory degradation. Finally, it is worth noting that, as it is the case for any other MVF estimation technique, the proposed method cannot distinguish between a production noise and a background noise. As a consequence, if speech is recorded in noisy conditions, MVF might be underestimated compared to what has been actually produced by the speaker.

III. OBJECTIVE EVALUATION

This section aims at objectively comparing the proposed approach with P2V [5] and SLM [6], for which we used the authors' original implementations. The experimental protocol and the results are further described in the following sections.

A. Experimental Protocol

The database used for the objective evaluation consists of *semi-synthetic* signals obtained as follows. First, F_0 and 24 True-Envelope-based Cepstral Coefficients (TECC, [15]) were extracted from real recordings. F_0 was estimated using the Summation of the Residual Harmonics (SRH) algorithm [16]. TE is here required as standard cepstral analysis was reported in [12] to be inappropriate for singing voice analysis. TE was estimated using the COVAREP toolkit [17]. These files were further resynthesized using the DSM [8] by imposing a constant value of the MVF. 7 versions of each file were created, with a fixed MVF ranging from 1 kHz to 7 kHz by step of 1 kHz. This way of doing allows us to have a ground truth to compare our MVF estimates with, and therefore to perform an objective assessment. Audio recordings were taken from CMU ARTCIC [18] speech databases (with 2 male and 2 female voices) and from the LYRICS singing voice database [19] (with excerpts from 7 baritones, 3 counter-tenors and 3 sopranos). The whole database is made of 1400 files (across the 7 MVF values, and with a balance across speech and singing voice sounds) which were further split equally into a development and a test sets. The development set was used during the design of the proposed technique in Section II, primarily to estimate the parameters of the Gaussian distributions used in the ML criterion. The test set was used to derive the results presented in Section III-B.

To assess performance, we use the rate of estimates with an error falling below a certain threshold θ . By varying θ between 0 and an upper bound (here fixed to 1.5 kHz), an approximate Receiver Operating Characteristic (ROC) curve is obtained, which is usually summarized by a single metric: the Area Under the Curve (AUC). AUC values are then normalized such that an ideal method would yield an AUC of 1. This metric was preferred to the root mean squared error (RMSE), as RMSE would be excessively biased by gross estimation errors.

B. Results

The approximate ROC curve obtained for the singing voice excerpts is shown in Figure 3. The state-of-the-art techniques P2V and SLM are compared to 5 variants of the proposed method depending on whether the AS, IHPC and ICPC features are used alone or jointly. It can be first observed that the baseline techniques achieve a rather poor performance on singing voices, and that a great improvement is yielded with the proposed methods. Among these latter, the use of ICPC alone brings the lowest performance, while the best results are obtained by using AS and IHPC features jointly.

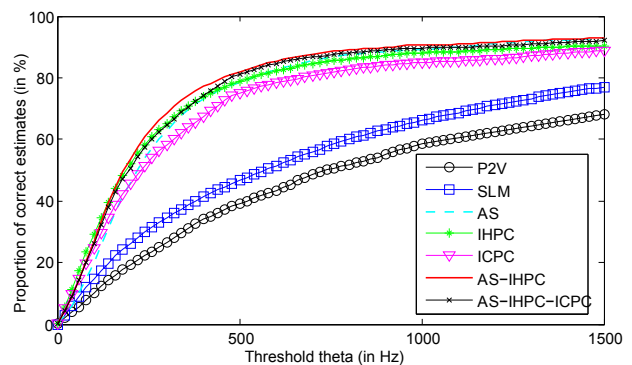


Fig. 3. Approximated ROC curve for the singing voice samples.

Figure 4 displays the AUC values averaged across all files for all compared methods. As it was observed in Figure 3, the proposed techniques clearly outperform state-of-the-art methods in singing voice. Interestingly, it can be noted that phase-based features achieve a rather high performance in both speech and singing voice, contrary to the use of the AS feature for which a degradation is noticed in speech. We also observed informally the advantage of phase-based features over AS to be more pronounced in female than in male speech. Among the proposed features, it turns out that IHPC carries out the best results, followed by ICPC and AS. According to these experiments, the combination of IHPC with other features does not seem to bring much improvement. Nonetheless, as it will be discussed in Section IV-A, we have observed on real signals that the joint use of several features tends to provide more stable and coherent estimates.

IV. SUBJECTIVE EVALUATION

Besides the objective evaluation on semi-synthetic signals, one might be interested in assessing the benefit of the proposed technique in the frame of voice synthesis. For this purpose we have carried out a subjective test whose details are given in Section IV-A and whose results are described in Section IV-B.




Fig. 4. Averaged AUC values for various MVF estimation techniques for speech and singing voice, together with their 95% confidence intervals.

A. Experimental Protocol

Again, our goal here is to cover a large diversity of voice signals. More precisely, six categories of voices are considered: male speech, female speech, child speech, baritones, counter-tenors, and sopranos. Adult speech samples are from AWB (male) and CLB (female) speakers of the CMU ARCTIC database [18]. Child speech sentences were from two 11 year-old female speakers, and were kindly provided by Acapela Group. In addition, samples of different singers were taken from the LYRICS database [19], for a total of 13 singers (7 bass-baritones, 3 countertenors, and 3 sopranos). The whole database used in our subjective evaluation contains in total 66 recordings mixed across the 6 categories. Note that none of these files were part of the datasets used in Section III.

By visual inspection of the MVF values estimated by the proposed approaches on these real signals, we noticed that combining various features led, in some cases, to more stable estimates. We also observed that the AS-IHPC and AS-IHPC-ICPC estimates were extremely close. However, AS and IHPC features require the computation of a single FFT per frame, while ICPC requires one more FFT (for the next pitch period) and therefore roughly doubles the computational cost. For these reasons, AS-IHPC was chosen as a representative of the proposed approach in the following experiments.

As in Section III-A, sounds were resynthesized using DSM with 3 feature streams as input: F_0 , 24 TECCs and the MVF estimated using either P2V, SLM or the proposed AS-IHPC technique. The evaluation consists of a Comparative Mean Opinion Score (CMOS) test, where participants are asked to compare two versions of the same file on a 7-point scale ranging from “much worse” (−3) to “much better” (+3). A nul score is given if both versions are found to be equivalent. In our experiments, two comparisons were considered: AS-IHPC vs. P2V, and AS-IHPC vs. SLM. Two groups of listeners took the test: 39 naive listeners (via the crowdsourcing website *CrowdFlower*) and 13 speech experts. Each participant was presented 20 randomized pairwise comparisons.

B. Results

Results of the CMOS test performed by the speech experts are exhibited in Figure 5. The general trend shows that the higher the pitch of the voice, the more the improvement of the proposed technique over the state of the art. Although the improvements in adult speech are in general not significant, AS-IHPC is clearly seen to outperform both P2V and SLM

in child speech and singing voice samples. This improvement even becomes substantial for counter-tenors and sopranos, for which averaged CMOS scores above 1.8 are obtained (which roughly implies that AS-IHPC is found to be “better”).




Fig. 5. Results of the CMOS test performed by the 13 speech experts. A positive score indicates a preference for the proposed AS-IHPC technique.

Results for naive listeners are presented in Fig. 6. Tendencies and conclusions drawn from these results are similar as with speech experts, except that they are in this case less pronounced. This can be explained by two facts: *i*) compared to speech experts, naive listeners pay less attention to details; *ii*) contrary to naive listeners, speech experts tend to use the full range of the CMOS scale. Nonetheless, it is appreciable to note that naive listeners also could observe differences between two versions of a voice sound which only differ by their estimated MVF contours.




Fig. 6. Results of the CMOS test performed by the 39 naive listeners. A positive score indicates a preference for the proposed AS-IHPC technique.

V. CONCLUSION

This paper proposed a new approach for MVF estimation. This technique was shown to outperform two state-of-the-art methods both in an objective and subjective evaluation. In the perceptual assessment, a substantial improvement was observed for high-pitched voices (particularly for counter-tenors and sopranos). This was reported for both speech experts and naive participants. The success of the proposed method lies in its three novelties: *i*) the use of the phase information, which is shown to have the potential to discriminate harmonicity; *ii*) the joint use of features extracted from the amplitude and phase spectra; *iii*) the use of a maximum likelihood criterion as a MVF decision strategy. A Matlab implementation of the proposed technique as well as audio examples can be found at tcts.fpms.ac.be/~drugman/Toolbox.

REFERENCES

- [1] T. Yoshimura, K. Tokuda, T. Masuko, T. Kitamura: *Mixed-excitation for HMM-based speech synthesis*, in Proceedings of Eurospeech, pp. 2259–2262, 2001.
- [2] H. Kawahara, J. Estill, O. Fujimura: *Aperiodicity extraction and control using mixed mode excitation and group delay manipulation for a high quality speech analysis, modification and synthesis system STRAIGHT*, in Proceedings of the 2nd International Workshop on Models and Analysis of Vocal Emissions for Biomedical Applications, 2001.
- [3] T. Raitio, A. Suni, H. Pulakka, M. Vainio, P. Alku: *Utilizing glottal source pulse library for generating improved excitation signal for HMM-based speech synthesis*, in Proceedings of ICASSP, pp. 4564–4567, 2011.
- [4] D. Griffin, J. Lim: *Multiband excitation vocoder*, in IEEE Trans. on Acoustics, Speech and Signal Processing, vol. 36, no. 8, pp. 1223–1235, 1988.
- [5] Y. Stylianou: *Applying the Harmonic plus Noise Model in Concatenative Speech Synthesis*, in IEEE Trans. on Speech and Audio Processing, vol. 9, pp. 21–29, 2001.
- [6] D. Erro, I. Sainz, E. Navas, I. Hernaez: *Harmonics plus Noise Model based Vocoder for Statistical Parametric Speech Synthesis*, in IEEE Journal of Selected Topics in Signal Processing, 2014.
- [7] Y. Pantazis, O. Rosec, Y. Stylianou: *On the Properties of a Time-Varying Quasi-Harmonic Model of Speech*, in Proceedings of Interspeech, pp. 1044–1047, 2008.
- [8] T. Drugman, T. Dutoit: *The Deterministic plus Stochastic Model of the Residual Signal and its Applications*, in IEEE Transactions on Audio, Speech and Language Processing, vol. 20, Issue 3, pp. 968–981, 2012.
- [9] T. Drugman, G. Wilfart, T. Dutoit: *A Deterministic plus Stochastic Model of the Residual Signal for Improved Parametric Speech Synthesis*, in Proceedings of Interspeech, pp. 1779–1782, 2009.
- [10] D. Erro, I. Sainz, I. Saratxaga, E. Navas, I. Hernaez: *MFCC+F0 Extraction and Waveform Reconstruction using HNM: Preliminary Results in an HMM-based Synthesizer*, in Proceedings of VI Jornadas en Tecnologia del Habla & II Iberian SLTech, pp. 29–32, 2010.
- [11] T. Drugman, T. Raitio: *Excitation Modeling for HMM-based Speech Synthesis: Breaking Down the Impact of Periodic and Aperiodic Components*, in Proceedings of ICASSP, 2014.
- [12] O. Babacan, T. Drugman, T. Raitio, D. Erro, T. Dutoit: *Parametric Representation for Singing Voice Synthesis: a Comparative Evaluation*, in Proceedings of ICASSP, 2014.
- [13] A. Ciobanu, T. Zorila, C. Negrescu, D. Stanomir: *Maximum Voiced Frequency Estimation for Voice Conversion used in text-to-Speech Systems*, in Proceedings of International Conference on Systems, Signals and Image Processing, pp. 476–479, 2012.
- [14] M. Zivanovic, A. Roebel, X. Rodet: *Adaptive Threshold Determination for Spectral Peak Classification*, in Computer Music Journal, vol. 32, number 2, pp. 57–67, 2008.
- [15] S. Imai: *Cepstral analysis synthesis on the mel frequency scale*, in Proceedings of ICASSP, vol. 8, pp. 93–96, 1983.
- [16] T. Drugman, A. Alwan: *Joint Robust Voicing Detection and Pitch Estimation Based on Residual Harmonics*, in Proceedings of Interspeech, pp. 1973–1976, 2011.
- [17] G. Degottex, J. Kane, T. Drugman, T. Raitio, S. Scherer: *COVAREP - A collaborative voice analysis repository for speech technologies*, in Proceedings of ICASSP, 2014.
- [18] J. Kominek, A. Black: *The CMU Arctic Speech Databases*, in Proceedings of the 5th Speech Synthesis Workshop, pp. 223–224, 2004.
- [19] N. Henrich, C. d’Alessandro, M. Castellengo, B. Doval: *Glottal open quotient in singing: Measurements and correlation with laryngeal mechanisms, vocal intensity, and fundamental frequency*, in JASA, vol. 117, pp. 1417–1430, 2005.