Performance Evaluation of Glottal Inverse Filtering Algorithms Using a Physiologically Based Articulatory Speech Synthesizer

Yu-Ren Chien, Daryush D. Mehta, *Member, IEEE,* Jón Guðnason, Matías Zañartu, *Member, IEEE,* and Thomas F. Quatieri, *Fellow, IEEE*

***Abstract*—Glottal inverse filtering aims to estimate the glottal airflow signal from a speech signal for applications such as speaker recognition and clinical assessment. Nonetheless, evaluation of inverse filtering performance has been challenging due to the practical difficulty in measuring the true glottal signals while speech signals are recorded. Apart from this, it is suspected that the performance of many methods degrade in conditions that are of great interest, such as breathy voice, high pitch, soft/loud voice, and running speech. This paper presents a comprehensive, objective, and comparative evaluation of state-of-the-art inverse filtering algorithms that takes advantage of speech and glottal signals generated by a physiologically relevant speech synthesizer. The synthesizer provides a realistic simulation of the voice production process, and thus an adequate test bed for revealing the temporal and spectral performance characteristics of each algorithm. Included in the synthetic data are continuous running speech utterances and sustained vowels, which are produced with multiple voice qualities (pressed, slightly pressed, modal, slightly breathy, and breathy) and subglottal pressure levels to simulate the natural variations in real speech. In evaluating the accuracy of a glottal flow estimate, multiple error measures are used, including an error in the estimated signal that measures overall waveform deviation, as well as an error in each of several clinically relevant features extracted from the glottal flow estimate. For two vowel-specific data subsets that were isolated for two open vowels and analyzed with three closed-phase approaches, the resulting waveform errors had mean and standard deviation values below 20% and 10%, respectively, of the true glottal source amplitude. These approaches also showed remarkable stability across different voice qualities and subglottal pressure levels. Results of data subset analysis suggest that analysis of close rounded vowels is a major challenge in glottal flow estimation.**

Distribution A: approved for public release; unlimited distribution.
This work is sponsored in part by The Icelandic Centre for Research (RANNIS) under the project *Model-based speech production analysis and voice quality assessment*, Grant No 152705-051, in part by the NIH National Institute on Deafness and Other Communication Disorders under Grant R21/R33 DC011588, and in part by CONICYT grants BASAL FB0008 and FONDECYT 1151077. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the NIH. This material is based upon work supported by the Assistant Secretary of Defense for Research and Engineering under Air Force Contract No. FA8721-05-C-0002 and/or FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Assistant Secretary of Defense for Research and Engineering.

Y.-R. Chien and J. Guðnason are with Center for Analysis and Design of Intelligent Agents, Reykjavik University, Menntavegur 1, Iceland (e-mail: yrchien@ntu.edu.tw, jg@ru.is).

D. D. Mehta is with the Center for Laryngeal Surgery and Voice Rehabilitation, and Institute of Health Professions, Massachusetts General Hospital, Boston MA 02114 USA, with the Department of Surgery, Harvard Medical School, Boston, MA 02115 USA, and also with MIT Lincoln Laboratory, Lexington, MA (e-mail: mehta.daryush@mgh.harvard.edu).

M. Zañartu is with the Department of Electronic Engineering, Universidad Técnica Federico Santa María, Valparaiso, Chile, 2390123 (e-mail: matias.zanartu@usm.cl).

T. F. Quatieri is with MIT Lincoln Laboratory, Lexington, MA (e-mail: quatieri@ll.mit.edu).

***Index Terms*—Performance evaluation, glottal flow estimation, inverse filtering, speech synthesis, glottal excitation, voice production, speech analysis.**

I. Introduction

HUMAN voice is the result of the glottal airflow exciting the vocal tract to produce the airflow through the lips and nostrils. Since the glottal airflow is modulated by the diaphragm and the vocal folds, which are in turn coordinated by the brain through motor control, an accurate estimate of the glottal excitation from a speech signal may provide salient information related to the speaker's identity, vocal function, emotions, etc. This makes glottal flow estimation desirable for speaker identification [1], voice quality assessment [2], analysis of emotional and neurological disorders [3], and clinical voice assessment [4], [5]. Nevertheless, true glottal signals have been elusive not only in ecological applications, but also in experimental settings. As a result, it has been difficult for researchers to evaluate the performance of a glottal flow estimator with confidence.

This paper presents an evaluation for a special class of glottal flow estimation methods, which we refer to as *inverse filtering* algorithms. As with manual inverse filtering procedures [6], [7],[1] these algorithms do not constrain the waveform estimate with a full-cycle glottal flow model; rather, less constrained glottal-flow assumptions are made where the inverse of vocal tract filter is estimated and applied to the speech signal to give a glottal flow estimate. Owing to this, inverse filtering algorithms are free from a performance limitation resulting from any deviation of real glottal flow characteristics from a glottal flow model, provided that an optional glottal flow modeling procedure following inverse filtering (such as the one presented in [1]) is not performed. In addition, for the glottal flow estimation techniques that are based on a glottal flow model (and thus not considered to be inverse filtering algorithms), the objective is typically to estimate only a subset of all the parameters required for glottal flow reconstruction, leaving a glottal signal estimate not well-defined. Consequently, among all the existing approaches to glottal flow estimation, only inverse filtering algorithms

[1]In a typical procedure for manual inverse filtering, an inverse filter (with user-specified formant frequencies and formant bandwidths) is manually adjusted and applied to an oral airflow signal from a flow mask, so that the output signal is free from ripples in the closed phase and has a smooth spectrum envelope, thus giving an estimate of the glottal airflow.

are tested in this study. In the evaluation, we aim to use synthesized glottal signals as a reference, test inverse filtering algorithms on speech signals synthesized from these glottal signals, and produce an objective assessment on the overall accuracy of each glottal waveform estimate.

In the experiments presented in this paper, both continuous running speech and sustained vowels are used for performance evaluation. The specific synthesis procedures adopted to generate these test materials are physiologically based, not only simulating the voice production mechanisms at the vocal fold and vocal tract levels, but also providing the ground-truth glottal source signals needed for the evaluation as part of the simulation. For sustained vowels, the data set includes synthesized speech utterances for various voice qualities and subglottal pressure levels. The resulting glottal source estimates are compared to the simulated glottal signals by measuring errors in time sample values, as well as in several types of feature values extracted from the waveform. Moreover, for the inverse filtering algorithms that make use of glottal closure instants detected from the speech signal, we evaluate the robustness to the errors in glottal closure detection with a simulation, where glottal closure instants are extracted from the synthesized glottal source signals, perturbed with controlled errors, and used to test these algorithms.

Our contribution is presented in the subsequent sections as follows. Related works are surveyed in Section II. The tested algorithms are reviewed in Section III. In Section IV, details are provided on how the sustained-vowel and continuous-running-speech data sets are constructed, and the performance measures used in the evaluation are also described. In Section V, results of our glottal flow estimation experiments are documented and analyzed for the tested algorithms. These results include examples that illustrate the ground-truth and estimated glottal source signals, as well as performance statistics calculated at the data-set level. Concluding remarks are given in Section VI.

II. Background

In a number of glottal flow estimation methods that do not perform inverse filtering, the parameters of a glottal flow model are estimated jointly with the parameters of a vocal tract filter. In an algorithm presented by Ding *et al.* [8], parameters are estimated from speech waveforms for the Rosenberg-Klatt (RK) glottal source model and a time-varying pole-zero-filter vocal tract model, by Kalman filtering and simulated annealing. Lu and Smith [9] estimated parameters of the KLGLOTT88 glottal source model and an all-pole vocal tract filter by solving a convex optimization problem that depends on detected glottal closure instants. In an analysis method presented by Funaki *et al.* [10], several models are adopted, including the RK glottal source model, a white-Gaussian random process for the aspiration noise, and a time-varying pole-zero filter for the vocal tract. They used the genetic algorithm as well as the technique of simulated annealing to fit these models to a speech signal, with phase distortion compensated by an all-pass filter. Fröhlich *et al.* [11] estimated parameters of an exponential-trigonometric (Liljencrants-Fant) glottal flow derivative model with a modified discrete all-pole modeling technique that optimizes the quality of inverse filtering. Vincent *et al.* [12] used the Liljencrants-Fant model and a time-varying all-pole-filter model for the vocal tract, with some parameters prioritized in a low-frequency analysis. Degottex *et al.* [13] used a minimum-phase vocal tract model to estimate the shape parameter and time position of the transformed Liljencrants-Fant model, and evaluated the resulting estimates with a digital vocal tract simulator. Model-based glottal flow estimation can also be achieved by fitting a glottal flow model to the glottal flow estimate given by an inverse filtering algorithm, as presented by Plumpe *et al.* [1].

In the case of inverse filtering algorithms, the glottal flow is defined with a representation more general than a parameterized waveform. Alku [14] presented a method for glottal excitation estimation that is based on representing the glottal source with a low-order linear-predictive spectrum envelope. Wong *et al.* [15] conducted linear-predictive covariance analysis in the closed phase of glottal pulse, and showed that the analysis implements least-squares estimation of the vocal tract filter, and that the closed phase can be located with a normalized error energy. Alku *et al.* [16] performed a closed-phase analysis where the inverse filter is constrained in terms of DC gain and minimum phase. They carried out performance evaluation with the vowel /a/ synthesized by a physical model of voice production that allows for simulation of the interaction between glottal source and vocal tract. To achieve better robustness to the errors in closed phase detection, Airaksinen *et al.* [17] estimated the vocal tract from both closed- and open-phase time samples with more weight on the closed-phase samples, and also evaluated their approach with physical modeling. Airaksinen *et al.* [18] recently modified the traditional closed-phase analysis by introducing an additional 1-norm term in the objective function of linear prediction. Drugman *et al.* [19] assumed a maximum-phase glottal open phase and minimum-phase signals for the glottal return phase and the vocal tract impulse response, and were able to estimate the open-phase glottal signal by causal-anticausal separation in the complex-cepstrum domain. In a different but related approach, Zañartu *et al.* [20] presented a non-parametric scheme to remove subglottal resonances in order to obtain glottal airflow estimates from a neck surface accelerometer. This case differs from the others in that it was designed for a different sensor and sensing position, and thus is not considered in this evaluation.

Glottal flow estimation is an important task in speech analysis for which performance evaluation and literature survey have been conducted in some dedicated works. Alku [21] reviewed the literature in the topics of glottal inverse filtering, parameterization of glottal flow estimates, and applications of inverse filtering, thereby concluding that the main current limitations of most inverse filtering methods are in high-pitch, running-speech, and pathological scenarios. Drugman *et al.* [22] evaluated three inverse filtering algorithms on real speech data with voice quality labels, as well as on synthetic speech data. Chu *et al.* [23] tested two closely-related inverse filtering algorithms with a sound-producing instrument modeled after the glottis and vocal tract. Drugman *et al.*

[24] presented a review of works on the glottal processing of speech, covering the aspects of synchronization, estimation, parameterization, and applications. More recently, Guðnason *et al.* [25] evaluated the performance of five inverse filtering algorithms with sustained vowels generated by an articulatory speech synthesizer, VocalTractLab [26].

In many of the above-mentioned works, glottal flow estimation experiments were conducted on synthetic audio data that is based on a shape-descriptive glottal flow model and an autoregressive vocal tract filter. Indeed, simplifications involved in such a model of voice production can result in inadequate synthesis, which in turn can give rise to a substantial performance gap between synthesized speech and real speech. This performance gap is especially relevant when many analysis approaches are actually based on the same models as the typical data synthesis procedure. In view of this, a small number of studies have drawn on physical modeling (either with numerical methods [13], [16], [17], [25], [27], [28] or with physical materials [23]) to fulfill realistic simulations of sustained vowels for the evaluation. In this work, we take a further step in enhancing the reality of test speech materials, by generating test data with the physiologically relevant synthesizer VocalTractLab, which is capable of synthesizing connected utterances by simulating user-specified articulatory movements. Furthermore, this study also expands on [25] by including multiple voice qualities and subglottal pressure levels in the test data, adopting several feature-based measures in performance evaluation, and performing a robustness analysis with respect to the errors in glottal closure detection.

III. Tested Algorithms

Five inverse filtering algorithms are tested in this evaluation:[2]

- closed-phase covariance analysis (CPCA) [15], which performs the estimation of vocal tract filter within the closed phase of glottal flow;
- sparse linear prediction (SLP) [29], which utilizes a (Gaussian-based) soft, weighted model of the glottal closed phase for a reduced dependency on the accurate detection of glottal closure and opening instants;
- weighted linear prediction (WLP) [30], which utilizes another (piecewise-linear) soft, weighted closed-phase model for the same purpose as with SLP;
- iterative adaptive inverse filtering (IAIF) [14], which is the most widely applied algorithm that estimates the glottal flow without assuming any glottal flow model or glottal closure instants; and
- complex cepstrum decomposition (CCD) [19], which is orthogonal to the other four algorithms in the sense that it is based on complex-cepstrum processing instead of linear prediction.

The descriptions in this section are specific to a custom implementation of each algorithm, except that Drugman's official implementation is used for CCD.[3] All the algorithms operate at the sampling frequency of 20 kHz in our implementation, with all synthesized signals resampled from their original sampling frequency of 44.1 kHz. Each algorithm is applied to a uniformly spaced sequence of time frames in the analyzed utterance. Since no glottal cycle exists within any non-voiced time interval in the utterance, any glottal signal estimated at a non-voiced time frame will be ignored by a cycle-synchronous measure when the accuracy of glottal flow estimation is evaluated at the utterance level.

A. Closed Phase Covariance Analysis (CPCA)

At each analysis time position,[4] say the τth position $n = n_\tau$, the vocal tract filter can be estimated by a linear-predictive covariance analysis that minimizes residual energy at closed-phase time samples [15]. Let the speech signal be denoted by $s[n]$, and let the vocal tract filter take the following form:

$$V(z) = \frac{1}{1 + \sum_{k=1}^{p} a_k z^{-k}}, \tag{1}$$

where p is set to 20.[5] The analysis calculates

$$a_k = -([b_{i,j}]^{-1}_{N\times(p+1)}[c_i]_{N\times 1})_{k+1},\ k = 1, ..., p, \tag{2}$$

where $(\cdot)_{k+1}$ denotes the $(k+1)$th element of a vector, and N is the window length (32 ms).[6] The matrix $[b_{i,j}]_{N\times(p+1)}$ is defined by

$$b_{i,j} = \begin{cases} w[n_\tau + i - 1], & \text{if } j = 1; \\ s[n_\tau + i - j]w[n_\tau + i - 1], & \text{otherwise}, \end{cases} \tag{3}$$

$$(i, j) \in \{1, ..., N\} \times \{1, ..., p + 1\}, \tag{4}$$

where $w[n]$ is unity if n is within the closed phase or within a non-voiced time interval, otherwise assuming the value zero. The vector $[c_i]_{N\times 1}$ is defined by

$$c_i = s[n_\tau + i - 1]w[n_\tau + i - 1],\ i = 1, ..., N. \tag{5}$$

Closed-phase boundaries are estimated by the YAGA algorithm [31]. The glottal flow derivative $\hat{\epsilon}[n]$ is then estimated by applying the inverse filter to the speech signal:

$$\hat{\epsilon}[n] = s[n] + \sum_{k=1}^{p} a_k s[n - k],\ n = n_\tau, ..., n_\tau + N - 1. \tag{6}$$

[2]For the reasons given in Section I, other approaches to glottal flow estimation are not tested in this study.

[3]http://tcts.fpms.ac.be/~drugman/Toolbox/GLOAT.zip. Customization of this implementation has been considered, but a major performance issue with this implementation has not been observed that could be resolved by a customization.

[4]For CPCA, SLP, and WLP, analysis time positions are spaced with a hop size of 16 ms. The hop size used in [15] was unspecified.

[5]In [15], the setting for the filter order was $p = 8$, with the sampling frequency unspecified. Since a filter of order 8 is typically used to model 4 formants for frequencies below 4 kHz, the sampling frequency there could have been 8 kHz.

[6]The window length used in [15] was 4.75 ms if a sampling frequency of 8 kHz was used. This ensured a time resolution that was sufficiently high for identifying the closed phase from linear-predictive residuals.

B. Sparse Linear Prediction (SLP)

As with CPCA, SLP estimates the vocal tract filter by a linear-predictive covariance analysis. However, this analysis minimizes a weighted sum of residual energy at all the time samples, with higher weights allocated to time samples farther from glottal closure instants [29]. Also using the equations (2), (3), and (5), the analysis defines its own weighting as follows:

$$w[n] = 1 - \kappa \cdot \sum_{l=1}^{L} \exp \frac{-(n-\gamma_l)^2}{2(\sigma f_s)^2}, \tag{7}$$

where γ_l denotes the lth of a total of L glottal closure instants detected from the speech signal [31], f_s denotes the sampling frequency in Hz, and κ and σ are parameters fixed to predefined constants (0.9 and 0.25 ms, respectively).[7]

C. Weighted Linear Prediction (WLP)

The WLP algorithm differs from SLP only in that its weighting is defined by a piecewise-linear function [30], rather than by a sum of upside-down, shifted Gaussian functions. The weighting is characterized by two distinct levels of weight (1.0 and 0.05),[8] with the higher-level value taken by all the time samples that are at a distance from glottal closure instants. Shortly before each glottal closure instant, the weight begins to ramp down, reaching the lower-level value before the glottal closure instant.[9] After retaining the low value (for 0.4 times the fundamental period) past the glottal closure instant, the weight starts to ramp up (for 0.45 ms), going back to the higher level shortly after the glottal closure instant.

D. Iterative Adaptive Inverse Filtering (IAIF)

Prior to estimating the vocal tract filter, the spectral contribution of glottal flow derivative can be estimated and removed from the speech signal with a low-order linear predictive analysis [14]. IAIF is a two-pass procedure based on this concept. In the first pass, a first-order linear predictive autocorrelation analysis is applied to the speech signal to give an estimate of the glottal source spectrum envelope.[10] After applying an inverse filter of this envelope to the speech signal, a 20th-order linear predictive autocorrelation analysis is applied to the filtered signal to give an estimate of the vocal tract filter,[11] according to which a second inverse filtering procedure yields the estimated glottal flow derivative for the first pass. In the second pass, low-order (4th-order) linear predictive analysis is again used to estimate the source contribution, but applied to the glottal flow estimated in the first pass. Similarly to the first pass, two inverse filtering steps follow to give the final estimate of glottal flow derivative.

[7]The value of σ used in [29] was 4.42 ms. Note that glottal closure instants were detected in [29] by the algorithm of Drugman and Dutoit [32].

[8]The value used in [30] for the lower level of weight was 0.01, determined from a synthetic development data set with true glottal closure instants.

[9]Ramping down takes 0.45 ms, and the lower level is reached 0.32 times the fundamental period before the glottal closure instant.

[10]All the linear predictive analyses in IAIF are carried out with a window length of 32 ms and a hop size of 16 ms.

[11]In [14], the higher order of linear prediction was set to 10 with a sampling frequency of 8 kHz.

E. Complex Cepstrum Decomposition (CCD)

At each analysis time position, say the lth position $n = \gamma_l$ which coincides with the lth glottal closure instant detected from the speech signal,[12] the glottal flow can be estimated directly by separating a maximum-phase component from the speech signal, without first estimating a vocal tract filter [19]. The CCD algorithm approaches the separation by calculating the complex cepstrum of the speech signal:

$$\hat{\mathbf{x}} = \mathrm{DFT}^{-1}\{\log|\mathrm{DFT}\{\mathbf{x}\}| + j\angle\mathrm{DFT}\{\mathbf{x}\}\}, \tag{8}$$

where $\mathrm{DFT}\{\cdot\}$ denotes the discrete Fourier transform, $\angle(\cdot)$ denotes the unwrapped phase of a complex number, and $\mathbf{x}$ denotes a time frame of the speech signal $s[n]$ centered at $n = \gamma_l$, spanning 1.8 cycles, multiplied by a Blackman window, and zero-padded to 102.4 ms.[13] The maximum-phase component is represented by the anti-causal component $\hat{\mathbf{x}}'$ in the complex cepstrum:

$$(\hat{\mathbf{x}}')_i = \begin{cases} \frac{1}{2}(\hat{\mathbf{x}})_1, & \text{if } i = 1; \\ 0, & \text{if } 2 \le i \le N_0/2; \\ (\hat{\mathbf{x}})_i, & \text{if } N_0/2 < i \le N_0, \end{cases} \tag{9}$$

where N_0 denotes the length of $\mathbf{x}$. The time-domain representation of the glottal flow estimate is then given by inverting the complex-cepstral calculation:

$$\mathbf{x}' = \mathrm{DFT}^{-1}\{\exp(\mathrm{DFT}\{\hat{\mathbf{x}}'\})\}, \tag{10}$$

from which an estimate of the glottal flow derivative can be calculated by taking the differences between adjacent elements.

IV. Experimental Procedure

A. Data Sets

All the utterances used in our experiments are generated by the software VocalTractLab 2.1 [26]. The synthesis of vowels performs time-domain, finite-difference simulation of acoustic wave motion for a two-mass, triangular-glottis model of the vocal folds [33] and a transmission-line model of the vocal tract. Since the triangular glottis model is time-varying, the synthesis is capable of simulating the time-varying coupling between glottal source and vocal tract [34]. Another physiological advantage of this glottis model is its capability of simulating a continuum of voice qualities from pressed voice to breathy voice. Voice quality concerns the degree of glottal closure within each glottal cycle, which can vary both within the same utterance and among different speakers. The pressed voice is characterized by a relatively long phase for closed vocal folds, whereas the vocal folds can lack a complete closure in the case of breathy voice. By being self-oscillating, the model promises more realistic glottal flow simulations than geometric approaches. The output sampling frequency of the synthesizer is 44.1 kHz. The approach is not currently capable of simulating pathological voices; therefore,

[12]In the implementation of CCD, the glottal closure instants $\gamma_1, ..., \gamma_L$ are detected by the algorithm of Drugman and Dutoit [32].

[13]As a default setting in Drugman's official implementation, the zero-padding length ensures a sufficiently high spectral resolution that is needed for phase unwrapping.

Table I
ERROR (MEAN ± STANDARD DEVIATION) OF GLOTTAL FLOW ESTIMATES ACROSS THE SUSTAINED-VOWEL AND CONTINUOUS-RUNNING-SPEECH DATA SETS. THE SUFFIX S REPRESENTS THE SIGNED VARIANT OF AN ERROR MEASURE. THE SUFFIXES LP AND HP REFER TO LOW- AND HIGH-PASS FILTERED VARIANTS OF MAEWAVE. THE ERROR GIVEN BY THE BEST-PERFORMING ALGORITHM IS SHOWN IN BOLDFACE FOR EACH COMBINATION OF DATA SET AND MEASURE. AS DEFINED IN SECTION IV-B, THE MEASURES MAEH1H2 AND MAEHRF (AND THEIR VARIANTS) ARE IN DB, AND THE OTHER MEASURES ARE UNIT-LESS.

Measure	Data	CPCA	SLP	WLP	IAIF	CCD
MAEWave	vowel	**0.27** ± 0.19	0.29 ± 0.18	0.29 ± 0.17	0.32 ± 0.20	0.33 ± 0.26
	speech	0.40 ± 0.11	**0.39** ± 0.12	0.39 ± 0.12	0.43 ± 0.11	0.43 ± 0.23
MAEWave-S	vowel	**0.00** ± 0.052	0.00 ± 0.062	0.001 ± 0.070	−0.008 ± 0.069	−0.074 ± 0.11
	speech	**−0.016** ± 0.026	−0.017 ± 0.028	−0.018 ± 0.027	−0.022 ± 0.027	−0.11 ± 0.13
MAEWave-LP	vowel	**0.24** ± 0.19	0.26 ± 0.18	0.26 ± 0.17	0.29 ± 0.20	0.33 ± 0.24
	speech	0.35 ± 0.097	**0.33** ± 0.10	0.34 ± 0.10	0.38 ± 0.10	0.43 ± 0.21
MAEWave-LP-S	vowel	0.014 ± 0.060	0.014 ± 0.064	0.012 ± 0.069	**0.002** ± 0.071	−0.13 ± 0.11
	speech	**−0.00** ± 0.030	−0.003 ± 0.033	−0.002 ± 0.031	−0.011 ± 0.031	−0.20 ± 0.14
MAEWave-HP	vowel	**0.088** ± 0.054	0.099 ± 0.055	0.099 ± 0.056	0.10 ± 0.057	0.099 ± 0.056
	speech	0.16 ± 0.089	0.16 ± 0.087	0.16 ± 0.089	0.16 ± 0.089	**0.15** ± 0.090
MAEWave-HP-S	vowel	**−0.00** ± 0.004	0.00 ± 0.005	0.00 ± 0.005	−0.00 ± 0.005	−0.002 ± 0.005
	speech	−0.00 ± 0.002	**0.00** ± 0.002	0.00 ± 0.002	−0.00 ± 0.002	−0.00 ± 0.002
MAENAQ	vowel	0.035 ± 0.027	0.031 ± 0.023	0.032 ± 0.024	**0.029** ± 0.023	0.044 ± 0.042
	speech	0.035 ± 0.017	0.034 ± 0.017	0.035 ± 0.016	**0.033** ± 0.015	0.042 ± 0.030
MAENAQ-S	vowel	0.030 ± 0.032	0.024 ± 0.029	0.026 ± 0.030	**0.020** ± 0.030	−0.042 ± 0.044
	speech	0.026 ± 0.026	0.026 ± 0.026	0.025 ± 0.027	**0.024** ± 0.024	−0.038 ± 0.032
MAEH1H2	vowel	3.3 ± 4.0	3.4 ± 4.1	**3.2** ± 3.9	4.2 ± 4.9	9.6 ± 7.2
	speech	3.1 ± 1.7	**3.0** ± 1.6	3.1 ± 1.7	3.6 ± 1.7	9.6 ± 5.1
MAEH1H2-S	vowel	**−0.91** ± 5.0	−2.5 ± 4.7	−2.1 ± 4.6	−3.2 ± 5.6	−9.4 ± 7.4
	speech	−0.57 ± 1.6	−0.51 ± 1.6	**−0.45** ± 1.5	−1.5 ± 1.6	−9.3 ± 5.3
MAEHRF	vowel	3.0 ± 2.9	3.0 ± 3.0	**2.7** ± 2.8	3.8 ± 4.3	12 ± 7.7
	speech	2.7 ± 1.6	**2.7** ± 1.6	2.7 ± 1.6	3.3 ± 1.7	13 ± 5.9
MAEHRF-S	vowel	**0.49** ± 4.1	2.2 ± 3.6	1.4 ± 3.6	2.9 ± 4.9	12 ± 7.7
	speech	0.45 ± 1.7	0.37 ± 1.7	**0.29** ± 1.7	1.6 ± 1.6	13 ± 6.0

we limit our analysis to the conditions currently included in VocalTractLab 2.1.

To evaluate the performance of inverse filtering algorithms under various controlled conditions, we carried out sustained-vowel time-domain simulation of voice production with VocalTractLab, giving a structured set of 750 speech utterances along with a corresponding set of glottal flow signals. These samples consisted of all the combinations of 5 target fundamental frequencies (for controlling vocal-fold tension; 90 Hz, 120 Hz, 150 Hz, 180 Hz, and 210 Hz), 5 subglottal pressure levels (500 Pa, 708 Pa, 1,000 Pa, 1,414 Pa, and 2,000 Pa), 5 voice qualities (pressed, slightly pressed, modal, slightly breathy, and breathy voice qualities), and 6 vowel types (/i/, /e/, /ɛ/, /ä/, /o/, and /u/). Each sample is a sustained-vowel utterance that lasts for 0.6 seconds.

A second data set is constructed for the continuous-running-speech experiments, which is generated by simulating manually planned movements in vocal-tract and vocal-fold configurations with VocalTractLab. All the utterances in this data set are derived from a prototype score of glottal and articulatory movements, which was composed by the author of VocalTractLab for the German sentence "Lea und Doreen mögen Bananen." The score describes 8 types of vocal movements, each of which is defined by a sequence of target configurations. Among the 8 movement types, three concern glottal movements,[14] i.e., target fundamental frequency (continuous-valued), subglottal pressure (continuous-valued), and voice quality (pressed, slightly pressed, modal, slightly breathy, or breathy). To generate utterances that exhibit different conditions of phonation, we adapted this prototype score by introducing various translations to the three glottal configuration sequences, such that each translated glottal configuration sequence has a new median value. The resulting adaptations consist of the 125 combinations of 5 median target fundamental frequencies, 5 median pressure levels, and 5 median voice qualities, which share specifications with the sustained-vowel data. The 125 new movement scores were used to synthesize 125 speech utterances, which make up our continuous-running-speech data set. In the adaptation, a translation by δ is introduced to the sequence of M voice quality values on the linear scale (with the 5 possible voice qualities encoded by the integers $1, ..., 5$):

$$\phi_m^{(\delta)} = \phi_m^{(0)} + \delta,\ m = 1, ..., M, \tag{11}$$

where $\phi_m^{(0)}$ and $\phi_m^{(\delta)}$ denote the mth prototype and translated voice quality values, respectively, such that the new sequence of voice quality values $\{\phi_m^{(\delta)}\}_{m=1}^{M}$ has one of the five desired median values while preserving the sequential variations in the prototype. Target fundamental frequencies (in Hz) and

[14]The other five types all concern vocal-tract movements.

subglottal pressures (in Pa) are similarly adapted, except that these are adapted on the logarithmic scale.

B. Performance Measures

Consider an utterance for which a glottal signal estimate has been produced by an inverse filtering algorithm. We assess the accuracy of the estimate in a cycle-synchronous fashion, accumulating cycle-wise error measurements over the whole utterance to give an overall measurement for the utterance. The utterance is segmented automatically into cycles according to the glottal area signals produced in the synthesis. To this end, the instantaneous minimum between the upper and lower glottal area signals is taken, and temporal variations in this minimum is monitored. The utterance is segmented whenever the minimum area drops below a threshold value that indicates glottal closure. Note that the instant when the glottal area vanishes does not always coincide with the glottal closure instant defined on the basis of glottal flow derivative [7]. The glottal area signals exhibit simple trends without impulse-like fluctuations, lending themselves to reliable detection of glottal closure events.

1) Waveform Errors: To determine the extent to which the estimated waveform deviates from the true glottal flow derivative, we calculate the *normalized median absolute waveform error* (MAEWave). Since the propagation delay in voice radiation is fixed in the synthesis and cannot be modeled by an inverse filtering algorithm in general,[15] the first step in assessing the accuracy of a glottal flow estimate is synchronization of the ground-truth glottal signal with the estimate. This is implemented by a 0.65-ms delay of the ground truth (corresponding to a 22-cm radiation distance), which gave the best alignment in our empirical experience with the synthesizer. Within a particular cycle, let the true and estimated glottal signals be denoted by $\epsilon_c[n]$ and $\hat{\epsilon}_c[n]$, respectively. Since signal amplitude is irrelevant to pulse shape comparison,[16] we calculate a scaled version of estimate whose amplitude is aligned with the true signal, with a scaling factor that minimizes the Euclidean distance between the scaled version and the true signal (i.e., by an orthogonal projection):

$$\tilde{\epsilon}_c[n] = \frac{\sum_{i=0}^{N_c-1} \epsilon_c[i]\hat{\epsilon}_c[i]}{\sum_{i=0}^{N_c-1} \hat{\epsilon}_c^2[i]} \cdot \hat{\epsilon}_c[n],\ n = 0, ..., N_c - 1, \quad (12)$$

where N_c denotes the length of this cycle. As shown in Fig. 1, a cycle-level waveform error is calculated by taking the error magnitude of $\tilde{\epsilon}_c[n]$ with respect to $\epsilon_c[n]$ for each time sample, taking the median of error magnitudes over all time samples in the cycle, and normalizing the median value by the utterance-wide root-mean-square amplitude of the true signal. The utterance-level waveform error is calculated by taking the median of all cycle-level errors.[17] Here the median-based measurement ensures that the resulting error accounts for a majority of its components, both on the cycle level and on the utterance level.

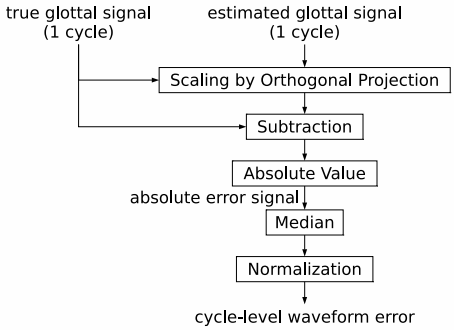

Figure 1. Waveform error evaluation for a particular cycle identified from the synthesized glottal area signals. The median of all cycle-level waveform errors in an utterance is calculated to give an MAEWave.

In the early days of voice production studies, inverse filtering used to be performed with dedicated hardware that came with no capability of optimization or matrix computation for formant frequency estimation, but allowed the user to assess glottal flow waveforms that resulted from various (user-specified) formant frequency settings [6]. The analysis implemented on a legacy inverse filtering device is typically limited to a bandwidth that only accounts for the first formant of vocal-tract frequency response. In the present study, to evaluate the accuracy of an estimated waveform in terms of what would have been given by single-formant processing, a variant of the aforementioned waveform error is calculated by applying the same error evaluation procedure to a low-pass filtered version of the true signal and a low-pass filtered version of the estimated signal. The low-pass filter is a 10th-order digital Butterworth filter with a cut-off frequency of 1 kHz [6]. On the other hand, to measure the higher-formant error component that could not be observed from single-formant processing, another variant of MAEWave is similarly calculated with a high-pass filter cut off at 1 kHz.

2) Feature Errors: The accuracy of a glottal flow estimate can also be assessed in terms of important waveform features that traditionally represent voice quality. To that end, we use the normalized amplitude quotient (NAQ) [35], the H1-H2 feature [36], and the harmonic richness factor (HRF) [37], calculating the *median absolute NAQ, H1-H2, and HRF errors* (MAENAQ, MAEH1H2, and MAEHRF). For each cycle of the true signal $\epsilon_c[n]$, an NAQ is evaluated by dividing the peak-to-peak glottal flow amplitude by the product of fundamental period and maximum flow declination rate. The maximum flow declination rate refers to the maximum magnitude of negative slope on the pulse shape of glottal flow (i.e., magnitude of the lowest point in the derivative pulse shape), which apparently

[15]Since the inverse filter is supposed to cancel both the magnitude and phase of the frequency response of vocal tract filter, neither the group delay of vocal tract filter nor that of the inverse filter is expected to contribute to a time delay between the estimated and ground-truth glottal flow signals.

[16]Since the relation between glottal pulse shape and voice quality is well-known, we focus on pulse shape comparison in this paper despite the potential importance of glottal signal amplitude in clinical applications.

[17]The utterance-level waveform error is not equivalent to a median calculated over all time samples in an utterance because the number of time samples within each cycle can vary from one cycle to another.

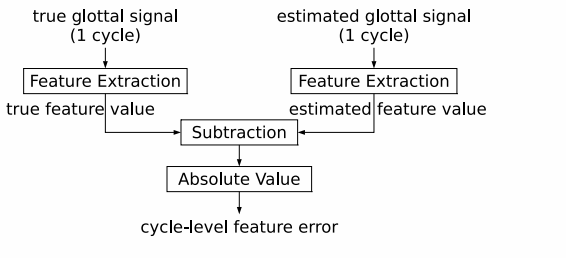

Figure 2. Error evaluation for a particular cycle (identified from the synthesized glottal area signals) and each of the features NAQ, H1-H2, and HRF. For each feature, the median of all cycle-level errors in an utterance is calculated to give a median absolute feature error.

varies with the fundamental frequency and signal amplitude. The NAQ feature eliminates this variance by normalizing the maximum rate by the diagonal slope of the rectangle spanned by the single-cycle pulse shape of the glottal flow. The features H1-H2 and HRF are also extracted from the true glottal flow as spectral descriptors of the single-cycle pulse shape. H1-H2 subtracts the amplitude of the second harmonic (in decibels) from the amplitude of the first harmonic. HRF measures the total power (in decibels) of overtones, i.e., the harmonics with an order greater than one, relative to the power of the fundamental. Here the harmonic amplitudes of the true glottal flow (integral of $\epsilon_c[n]$) are calculated by taking the absolute value of its discrete Fourier transform (without zero-padding before the transform) and extracting the consecutive frequency bins that correspond to harmonic frequencies greater than 0 Hz and less than 3 kHz. Both NAQ and H1-H2 could be regarded as a measure of voice breathiness, while HRF is negatively correlated with breathiness [35]–[37]. The same features are also extracted from the glottal signal estimate $\hat{\epsilon}_c[n]$. As shown in Fig. 2, to evaluate the error in glottal flow estimation, three error magnitudes are calculated respectively for the three features for each cycle, and an utterance-level error is calculated by taking the median of all cycle-level errors for each of the three features.

C. Simulation of Glottal Closure Instants

To evaluate the susceptibility of inverse filtering algorithms to the errors in glottal closure detection, we extract all the glottal closure instants from each true glottal flow signal in the data set, use these true instants to simulate estimated instants of a certain accuracy, and substitute these simulated estimates for the real detector-produced estimates in a glottal flow estimation experiment.

To extract glottal closure instants from a true glottal signal, the signal is first segmented into cycles with the same area-based procedure as described in Section IV-B. In each cycle, the time sample where one finds the lowest point of derivative pulse shape is extracted as a true glottal closure instant.

The error in an estimated glottal closure instant can be measured in relation to the instantaneous fundamental period, as a phase error in the quasi-periodic structure of glottal closure instants. To see the effect that this type of phase errors have on the performance of glottal flow estimation, we simulate estimates of glottal closure instants that have a constant phase error of θ radians throughout an utterance:

$$\tilde{\gamma}_l = \left\lfloor \bar{\gamma}_l + (\bar{\gamma}_{l+1} - \bar{\gamma}_l) \cdot \frac{\theta}{2\pi} + 0.00065 f_s + 0.5 \right\rfloor, \quad (13)$$

$$l = 1, ..., L-1, \quad (14)$$

$$\tilde{\gamma}_L = \left\lfloor \bar{\gamma}_L + (\bar{\gamma}_L - \bar{\gamma}_{L-1}) \cdot \frac{\theta}{2\pi} + 0.00065 f_s + 0.5 \right\rfloor, \quad (15)$$

where $\tilde{\gamma}_l$ denotes the lth simulated glottal closure instant in samples, $\bar{\gamma}_l$ denotes the lth true glottal closure instant in samples, f_s denotes the sampling frequency in Hz, and a rounding to the nearest integer and a 0.65-ms delay give the simulated estimate.

V. Results

A. Results on Sustained Vowel and Continuous Running Speech Data

Results of the sustained-vowel and continuous-running-speech glottal flow estimation experiments are presented in Table I. For sustained vowels, all the five algorithms gave normalized waveform errors around 0.3, with standard deviations around 0.2, which shows no substantial performance difference among the algorithms. Listed on the row titled "MAEWave-S" are results obtained with a signed variant of the waveform error, where a signed error is calculated in place of an error magnitude for each time sample to reveal any systematic bias in the signal estimate. This shows that IAIF and CCD tend more to underestimate glottal flow derivative values than to overestimate them, but even for these two algorithms the bias does not predominantly account for the unsigned waveform error.

The similarity between the low-pass filtered and unfiltered waveform errors (measured as described in Section IV-B1) suggests a consistency of the present performance measurement with earlier research. Although large signal value errors could occur in the return phase (because of the typically abrupt change in glottal flow derivative) and thus be captured by the high-pass filtered measure, such errors would be confined within a small number of time samples in each cycle and have no substantial impact on the median-based high-pass measure. This explains why the low-pass error component dominates the waveform errors.

The NAQ results again show a similarity of performance among the algorithms, but reveal that errors in NAQ are overwhelmingly either underestimations (with a large, negative signed error for CCD) or overestimations (with a large, positive signed error for the other algorithms) within an algorithm. This suggests the possibility of improving NAQ estimates given by a specific algorithm by canceling the bias observed here. The results for the spectral features H1-H2 and HRF show relatively poor performance for CCD with average errors around 10 dB, and substantial biases (underestimations of H1-H2 and overestimations of HRF) for all the algorithms except CPCA.

Regarding the continuous-running-speech results, mean MAEWave was again similar (approximately 0.40) across all

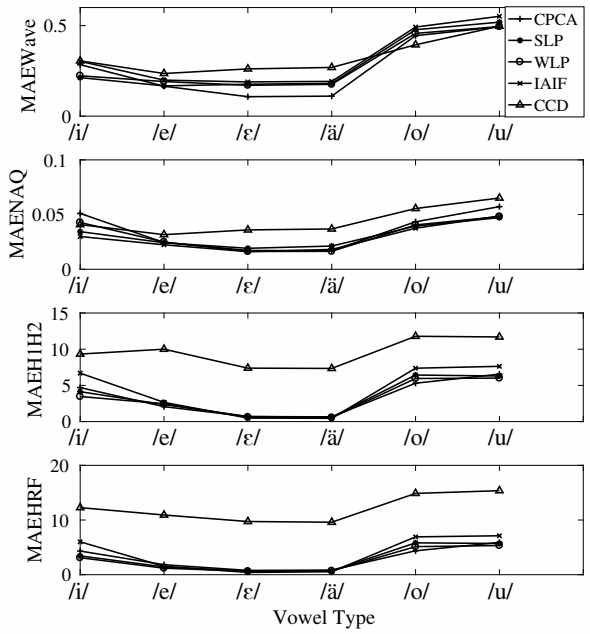

Figure 3. Subset error averages for vowel types.

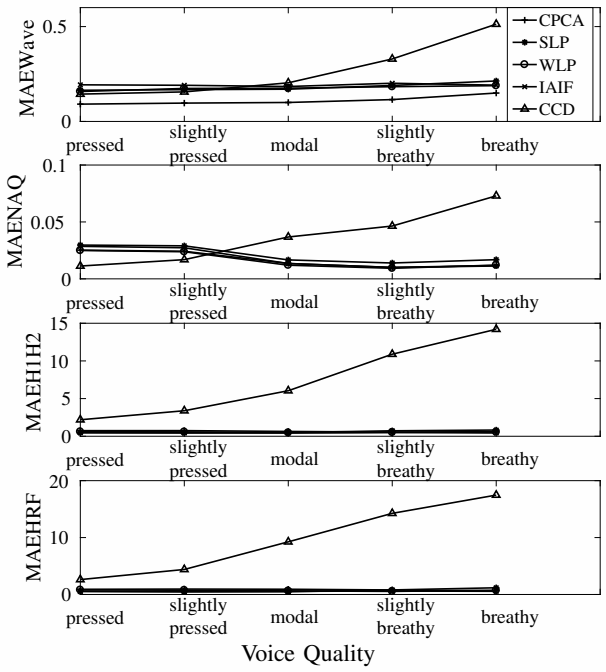

Figure 4. Subset error averages for voice qualities. Only utterances of vowel /ä/ in the sustained-vowel data set are used.

the algorithms, and comparison of the MAEWave results with those obtained with the variant measures exhibits an absence of substantial bias, as well as a consistency of unfiltered results with low-pass filtered results. The NAQ results reveal the biasedness of all NAQ estimates. CCD produced H1-H2 and HRF estimates with a bias that resulted in an average error on the order of 10 dB.

B. Results on Data Subsets

1) Vowel Types: It has been observed by some researchers that some vowels with a low first formant frequency cannot be adequately analyzed by an inverse filtering algorithm, whereas the vowel /ä/ has a first-formant frequency that is sufficiently high to avoid interference with the primarily low-frequency energy distribution of glottal source [14]. To see the impact of vowel type on the performance of algorithms, we took a separate average of errors for each vowel-specific subset of the sustained-vowel data. As shown in Fig. 3, the close rounded vowels /o/ and /u/ are associated with substantially higher errors than other vowels. This confirms that the analysis of close rounded vowels remains difficult as far as inverse filtering algorithms are concerned. Throughout the rest of this paper, we will move on to explore some other factors that could also have an effect on algorithm performance, while factoring out the effect of vowel types by testing the algorithms on utterances of the vowel /ä/ only.

2) Voice Qualities: The performance of algorithms on utterances of different voice qualities is examined in Fig. 4. For CCD, the breathy voice quality is associated with a drastically higher average error than the pressed voice quality with respect to every performance measure, which suggests that the maximum phase property assumed for the glottal open phase may not be as valid for breathy voice as for pressed voice. All the other algorithms demonstrate roughly constant performance over the voice qualities with respect to several measures. This is remarkable for the closed-phase approaches in particular, for which only a small number of time samples are available in each analysis time frame for the estimation of vocal tract filter in the case of breathy voice. An exception to this constant performance is the NAQ error, for which the pressed and slightly pressed voice qualities have slightly higher errors. This resulted from the narrow negative peaks in pressed glottal flow derivative waveforms, which are not represented accurately by the 20-kHz signal sampling in our experiments. Accurate performance evaluation in terms of the NAQ feature would require a sampling frequency higher than 44.1 kHz because even the un-resampled derivative waveforms from the synthesizer for pressed voice, exhibit maximum flow declination rates that vary substantially between adjacent cycles.

3) Subglottal Pressure Levels: Sustained-vowel utterances of a particular subglottal pressure are also isolated to give an average error specific to the pressure level. These errors are plotted in Fig. 5, where the only remarkable effect of the pressure level occurs with the waveform errors given by the CCD algorithm. The raised error for low pressure could be an effect similar to that of breathiness, for low subglottal pressure tends to result in a glottal pulse shape typical of a breathy voice.

4) Target Fundamental Frequencies: Inverse filtering algorithms typically involve the estimation of vocal tract filter, which explicitly or implicitly relies on the harmonic amplitudes of speech signal as observable samples of the spectrum envelope. As the fundamental frequency increases, the observable harmonics become sparser in the spectrum, which can gradually turn the envelope estimation problem into an under-

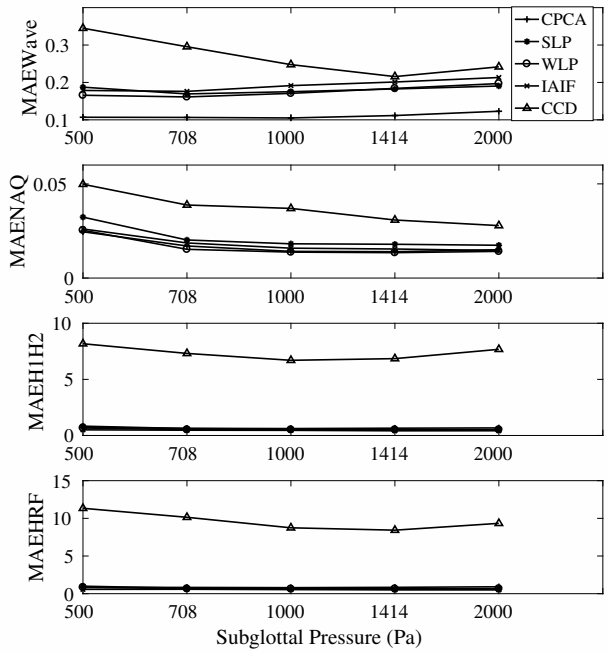

Figure 5. Subset error averages for subglottal pressure levels. Only utterances of vowel /ä/ in the sustained-vowel data set are used.

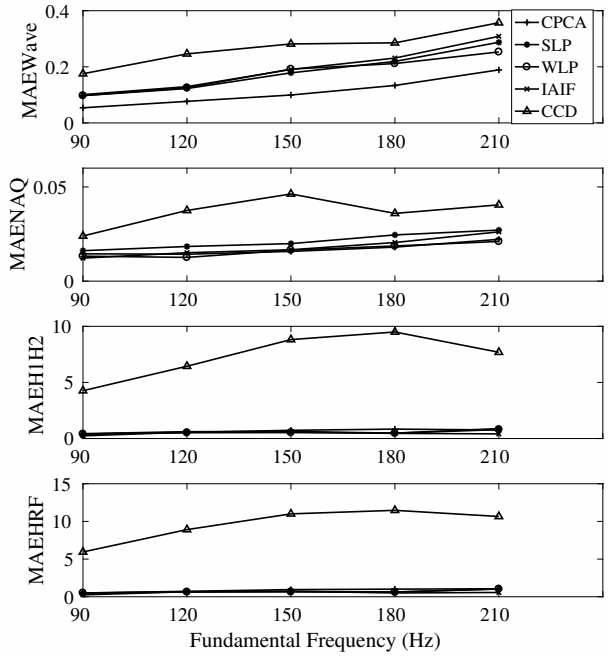

Figure 6. Subset error averages for target fundamental frequencies. Only utterances of vowel /ä/ in the sustained-vowel data set are used.

determined one. The degradation of glottal flow estimation performance under increasing fundamental frequency has been well documented and discussed in the literature, which is also observed on our data set as a general trend of MAEWave in Fig. 6. In comparison to the other algorithms, the evidently inferior performance of CCD presumably results from the limited validity of its assumption on the maximum-phase open-phase glottal flow, given that none of the others is based on the assumption.

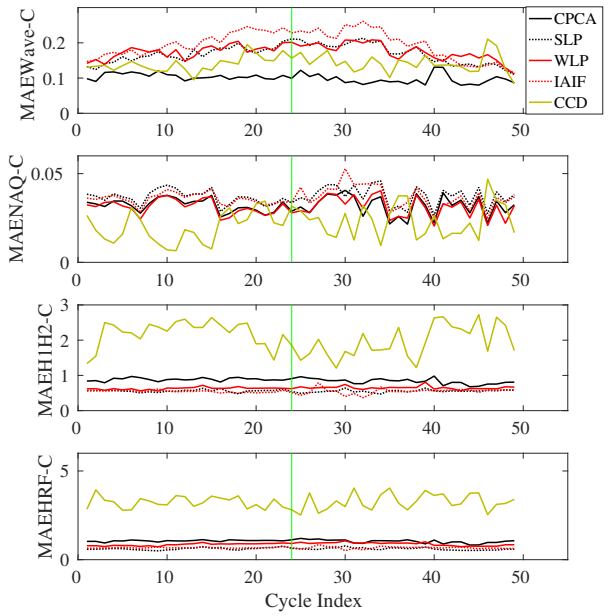

Figure 7. Cycle-level errors in a sustained-vowel example utterance, which has a slightly pressed glottal flow, vowel /ä/, a target fundamental frequency of 150 Hz, and a subglottal pressure of 500 Pa. The suffix C refers to the cycle-level errors underlying an utterance-level measure. Marked with a vertical green line is the cycle with the 25th lowest error among the 49 cycles in terms of the CPCA algorithm and the cycle-level components of MAEWave.

C. Examples

To give an example that is representative of the glottal flow estimation performance, we use a specific utterance for which median performance was observed among all the utterances in the data set. The median-performance utterance is determined in terms of the MAEWave measure and the CPCA algorithm. The utterance is selected such that its error is the 63rd lowest (0.099) among all the 125 utterances of vowel /ä/. For this example utterance, results can be examined not only in terms of MAEWave and CPCA, but also in terms of other measures and algorithms. Cycle-level errors are plotted for this utterance in Fig. 7, which shows that the best-performing algorithm varies on the utterance level, depending on the error measure used: When the waveform error is used, CPCA gave the lowest error. When the NAQ error is used, the lowest error was given by CCD. When either of the two spectral-feature errors is used, IAIF and SLP performed the best.

Physiologically based speech synthesis could simulate a "ripple effect" in the glottal airflow that is beyond the representation of a typical glottal flow model. It consists in some ripples in the open-phase glottal flow derivative waveform that result from the nonlinear coupling between vocal tract and glottis [34]. To see how well these ripples can be captured by an inverse filtering algorithm, we assess the accuracy of glottal flow estimation also at the cycle level. To that end, we apply the same median selection strategy to the cycles in the example utterance, illustrating with the median-performance cycle determined in terms of the cycle-level components of

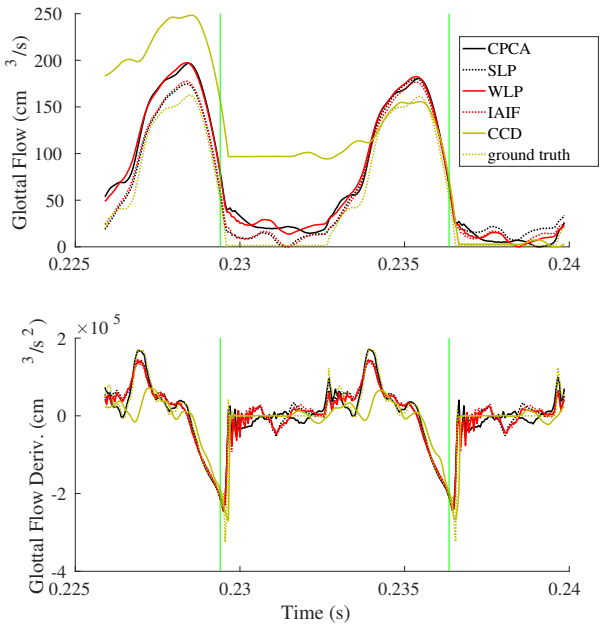

Figure 8. Glottal flow and derivative ($\tilde{\epsilon}_c[n]$) estimates generated by the tested algorithms at the median cycle marked in Fig. 7. The endpoints of the cycle are marked with vertical green lines.

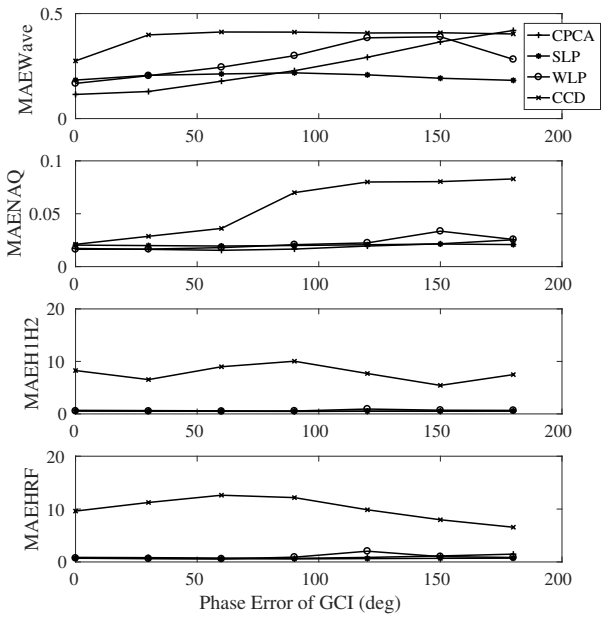

Figure 9. Performance of four algorithms under various amounts of error in each simulated glottal closure instant (GCI) estimate used by the algorithms. Only utterances of vowel /ä/ in the sustained-vowel data set are used in these experiments.

MAEWave and the CPCA algorithm. The example cycle is selected such that its error is the 25th lowest (0.099) among all the 49 cycles in the example utterance. The estimates given by the five algorithms for the selected cycle are shown in Fig. 8. In the derivative plot, CPCA slightly deviates from the ground truth during the closed phase, but closely matches the ground truth during the open phase, where the ripples are evident. In contrast, CCD closely matches the ground truth during the closed phase while deviating considerably during the open phase. The latter deviation is so severe that the algorithm received spectral-feature errors around or above 2 dB. Note that the glottal flow derivative estimate given here by CCD has a nonzero DC value, whose magnitude is so large that the corresponding glottal flow estimate (integral of the derivative) appears substantially aperiodic. All the glottal flow estimates show a maximum descending slope similar to the ground-truth slope, giving NAQ errors around 0.03.

D. Robustness to Errors in Glottal Closure Detection

Results for the simulated glottal closure detection is presented in Fig. 9. As intended by the weighted minimization of residual energy, the dependency of performance on the accuracy of glottal closure detection is minimum for SLP under every performance measure. The strong dependency for the other three algorithms is evident in terms of the waveform error. Note that IAIF does not rely on glottal closure detection. Since CPCA also uses glottal opening instants, the YAGA algorithm is used for generating candidates for the opening instants, which can be used along with the simulated glottal closure instants to give estimates for the opening instants.

VI. Conclusions

In this paper, the performance of several inverse filtering algorithms has been evaluated with synthesized test data. These algorithms aim to provide accurate glottal flow estimates without assuming a full-cycle glottal flow model. With the test data generated with a physiologically relevant, articulatory speech synthesizer that simulates articulatory movement as well as voice production, the resulting evaluation serves to predict the performance of these algorithms in analyzing real speech.

The fundamental techniques that underlie the tested methods include linear-predictive covariance analysis, linear-predictive autocorrelation analysis, and the complex cepstrum. The experiments showed that each method gives an average MAEWave around 0.3 over the sustained-vowel data, and an average error of the same type around 0.4 over the continuous-running-speech data. The average waveform errors evaluated over the close rounded vowel subsets of the sustained-vowel data are above 0.4 for all the methods, which confirmed that the methods are not as effective for close rounded vowels as for open vowels. Comparison among data subsets of an open vowel and of different voice qualities revealed that CCD does not produce glottal flow estimates as accurate for breathy voice as for pressed voice, which suggests that the validity of the maximum-phase assumption on open-phase glottal flow is questionable in the case of breathy voice. According to the robustness analysis performed with respect to the errors in glottal closure detection, the algorithm of choice for the analysis of vowel /ä/ is IAIF or SLP when accurate glottal closure instants are not available.

Results of the experiments suggest that the difficulty in analyzing close rounded vowels remains a major factor that limits

the applicability of inverse filtering algorithms to accurate glottal flow estimation from continuous running speech. This difficulty could have resulted from the first-formant resonance in close rounded vowels coinciding with the frequency band where glottal source energy is primarily distributed. It would be an important direction for future research to inquire models of voice production that are effective for the analysis of close rounded vowels. Other challenges in glottal flow estimation also merit further investigation, including high-pitched phonation, disordered speech, and estimation from non-audio signals such as oral airflow and neck-surface accelerometry.

ACKNOWLEDGMENT

The authors would like to thank Dr. Peter Birkholz for sharing a handcrafted score of vocal movements, and Dr. Mark R. P. Thomas for sharing his implementation of IAIF.

REFERENCES

[1] M. D. Plumpe, T. F. Quatieri, and D. A. Reynolds, "Modeling of the glottal flow derivative waveform with application to speaker identification," *IEEE Transactions on Speech and Audio Processing*, vol. 7, no. 5, pp. 569–586, Sep. 1999.

[2] D. H. Klatt and L. C. Klatt, "Analysis, synthesis, and perception of voice quality variations among female and male talkers," *J. Acoust. Soc. Am.*, vol. 87, no. 2, pp. 820–857, 1990.

[3] E. Moore, M. A. Clements, J. W. Peifer, and L. Weisser, "Critical analysis of the impact of glottal features in the classification of clinical depression in speech," *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 1, pp. 96–107, Jan. 2008.

[4] R. E. Hillman, E. B. Holmberg, J. S. Perkell, M. Walsh, and C. Vaughan, "Objective assessment of vocal hyperfunction: An experimental framework and initial results," *Journal of Speech, Language, and Hearing Research*, vol. 32, pp. 373–392, 1989.

[5] D. D. Mehta, J. H. Van Stan, M. Zañartu, M. Ghassemi, J. V. Guttag, V. M. Espinoza, J. P. Cortés, H. A. Cheyne II, and R. E. Hillman, "Using ambulatory voice monitoring to investigate common voice disorders: Research update," *Front Bioeng Biotechnol*, vol. 3, no. 155, 2015.

[6] J. S. Perkell, E. B. Holmberg, and R. E. Hillman, "A system for signal processing and data extraction from aerodynamic, acoustic, and electroglottographic signals in the study of voice production," *J. Acoust. Soc. Am.*, vol. 89, no. 4, pp. 1777–1781, 1991.

[7] S. Granqvist, S. Hertegård, H. Larsson, and J. Sundberg, "Simultaneous analysis of vocal fold vibration and transglottal airflow; exploring a new experimental set-up," *TMH-QPSR*, vol. 45, no. 1, pp. 35–46, 2003.

[8] W. Ding, H. Kasuya, and S. Adachi, "Simultaneous estimation of vocal tract and voice source parameters based on an ARX model," *IEICE Trans. Inf. & Syst.*, vol. E78-D, no. 6, pp. 738–743, 1995.

[9] H.-L. Lu and J. O. Smith III, "Joint estimation of vocal tract filter and glottal source waveform via convex optimization," in *Proc. 1999 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, 1999.

[10] K. Funaki, Y. Miyanaga, and K. Tochinai, "Recursive ARMAX speech analysis based on a glottal source model with phase compensation," *Signal Processing*, vol. 74, pp. 279–295, 1999.

[11] M. Fröhlich, D. Michaelis, and H. W. Strube, "SIM—simultaneous inverse filtering and matching of a glottal flow model for acoustic speech signals," *J. Acoust. Soc. Am.*, vol. 110, no. 1, pp. 479–488, 2001.

[12] D. Vincent, O. Rosec, and T. Chonavel, "Estimation of LF glottal source parameters based on an ARX model," in *Proc. INTERSPEECH*, 2005.

[13] G. Degottex, A. Roebel, and X. Rodet, "Phase minimization for glottal model estimation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 5, pp. 1080–1090, 2011.

[14] P. Alku, "Glottal wave analysis with pitch synchronous iterative adaptive inverse filtering," *Speech Communication*, vol. 11, pp. 109–118, 1992.

[15] D. Y. Wong, J. D. Markel, and A. H. Gray, "Least squares glottal inverse filtering from the acoustic speech waveform," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-27, no. 4, pp. 350–355, 1979.

[16] P. Alku, C. Magi, S. Yrttiaho, T. Bäckström, and B. Story, "Closed phase covariance analysis based on constrained linear prediction for glottal inverse filtering," *J. Acoust. Soc. Am.*, vol. 125, no. 5, pp. 3289–3305, 2009.

[17] M. Airaksinen, T. Raitio, B. Story, and P. Alku, "Quasi closed phase glottal inverse filtering analysis with weighted linear prediction," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 3, pp. 596–607, 2014.

[18] M. Airaksinen, T. Backstrom, and P. Alku, "Quadratic programming approach to glottal inverse filtering by joint norm-1 and norm-2 optimization," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, to appear.

[19] T. Drugman, B. Bozkurt, and T. Dutoit, "Causal-anticausal decomposition of speech using complex cepstrum for glottal source estimation," *Speech Communication*, vol. 53, pp. 855–866, 2011.

[20] M. Zañartu, J. C. Ho, D. D. Mehta, R. E. Hillman, and G. R. Wodicka, "Subglottal impedance-based inverse filtering of voiced sounds using neck surface acceleration," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 9, 2013.

[21] P. Alku, "Glottal inverse filtering analysis of human voice production—a review of estimation and parameterization methods of the glottal excitation and their applications," *Sādhanā*, vol. 36, no. 5, pp. 623–650, 2011.

[22] T. Drugman, B. Bozkurt, and T. Dutoit, "A comparative study of glottal source estimation techniques," *Computer Speech and Language*, vol. 26, pp. 20–34, 2012.

[23] D. T. W. Chu, K. Li, J. Epps, J. Smith, and J. Wolfe, "Experimental evaluation of inverse filtering using physical systems with known glottal flow and tract characteristics," *J. Acoust. Soc. Am.*, vol. 133, no. 5, pp. EL358–EL362, 2013.

[24] T. Drugman, P. Alku, A. Alwan, and B. Yegnanarayana, "Glottal source processing: From analysis to applications," *Computer Speech and Language*, vol. 28, pp. 1117–1138, 2014.

[25] J. Guðnason, D. D. Mehta, and T. F. Quatieri, "Evaluation of speech inverse filtering techniques using a physiologically based synthesizer," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015.

[26] P. Birkholz, "Modeling consonant-vowel coarticulation for articulatory speech synthesis," *PLoS ONE*, vol. 8, no. 4, 2013, doi:10.1371/journal.pone.0060603.

[27] P. Milenkovic, "Glottal inverse filtering by joint estimation of an AR system with a linear input model," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-34, no. 1, pp. 28–42, 1986.

[28] P. Alku, B. Story, and M. Airas, "Estimation of the voice source from speech pressure signals: Evaluation of an inverse filtering technique using physical modelling of voice production," *Folia Phoniatr Logop*, vol. 58, no. 2, pp. 102–113, 2006.

[29] V. Khanagha and K. Daoudi, "An efficient solution to sparse linear prediction analysis of speech," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2013, no. 3, 2013.

[30] P. Alku, J. Pohjalainen, M. Vainio, A.-M. Laukkanen, and B. H. Story, "Formant frequency estimation of high-pitched vowels using weighted linear prediction," *J. Acoust. Soc. Am.*, vol. 134, no. 2, pp. 1295–1313, 2013.

[31] M. R. P. Thomas, J. Gudnason, and P. A. Naylor, "Estimation of glottal closing and opening instants in voiced speech using the YAGA algorithm," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 1, pp. 82–91, 2012.

[32] T. Drugman and T. Dutoit, "Glottal closure and opening instant detection from speech signals," in *INTERSPEECH-2009*, 2009.

[33] P. Birkholz, B. J. Kröger, and C. Neuschaefer-Rube, "Synthesis of breathy, normal, and pressed phonation using a two-mass model with a triangular glottis," in *Interspeech 2011*, 2011.

[34] T. V. Ananthapadmanabha and G. Fant, "Calculation of true glottal flow and its components," *Speech Communication*, vol. 1, no. 3–4, pp. 167–184, 1982.

[35] P. Alku, T. Bäckström, and E. Vilkman, "Normalized amplitude quotient for parametrization of the glottal flow," *J. Acoust. Soc. Am.*, vol. 112, no. 2, pp. 701–710, 2002.

[36] I. R. Titze and J. Sundberg, "Vocal intensity in speakers and singers," *J. Acoust. Soc. Am.*, vol. 91, no. 5, pp. 2936–2946, 1992.

[37] D. G. Childers and C. K. Lee, "Vocal quality factors: Analysis, synthesis, and perception," *J. Acoust. Soc. Am.*, vol. 90, no. 5, pp. 2394–2410, 1991.