

A Quasi-glottogram signal.

Greg Kochanski and Chilin Shih

Bell Laboratories

Abstract

We propose a novel, noninvasive experiment that reliably shows the strength of glottal oscillations. The Quasi-glottogram (QGG) signal is generated from a microphone array that is trained to approximate the electroglottogram signal. The QGG may be useful to improve estimates of whether speech is voiced, to quantify partial voicing, and to reduce the phoneme effect when measuring the amplitude of speech signals. The technique is well adapted to the generation of text-to-speech systems, as it allows an estimate of the glottal flow during undisturbed, natural speech. For prosody studies, it can be used to provide an estimate of amplitude which is relatively unaffected by changes in phonemes, and is at least as reliable as standard estimators of amplitude.

PACS: 43.70.Jt, 43.70.Fq, 43.70.Aj, 43.58.Ta, 43.38.Hz

Abbreviated title: "A Quasi-glottogram signal"

Corresponding author: Greg Kochanski gpk@alum.mit.edu

1. Introduction

The source-filter model of speech production explains the acoustic speech signal as a convolution of the time-varying glottal airflow (due to the vibrations of the vocal folds) with the impulse response of the acoustic filter formed by the vocal tract. A measurement of the glottal airflow is clearly desirable: it would allow the source and the filter to be studied independently. In particular, it would allow a determination of voicing, and allow measurements of the amplitude and harmonic content of the glottal oscillations. The source-filter model of speech is embedded in speech coders, automatic speech recognition systems and speech synthesizers; better understanding of the source and filter separately can lead to better algorithms. For example, speech synthesizers need reliable, precise indications of glottal oscillation because voiced speech is often processed differently from unvoiced speech. Human listeners will easily detect errors in the database of a synthesizer where the speech has a voicing indication that is wrong for more than 30 ms, a sensitivity that requires improvement in current techniques if one wants to create synthesizers automatically without tedious manual checking.

Further, a relatively unexplored area of speech science is the study of prosody. We use the term “prosody” to include all the acoustic properties of speech that are not part of the lexical information. Prosody is typically implemented by properties such as the pitch, the amplitude, and the spectral tilt of the sound. Prosody is used to mark boundaries, to emphasize words or phrases, and to help control dialogs, among other functions. However, quantitative measurements of prosody are non-trivial, because other than pitch, all the candidate acoustic features are strongly influenced by the phonemes: for example, an emphatic ‘m’ may be quieter than a soft ‘a’. Thus, to be able to study and understand prosody, there is a need for measurements of prosody that are relatively independent of the particular phonemes, so that one can compare prosody of different words. Again, measurements of the glottal flow would be useful, because on one hand, the glottal flow is much less dependent on the choice of phoneme than the far-field acoustic signal is, and on the other hand because it can be related to physiological parameters like subglottal pressure and muscle tensions.

Invasive estimates of the glottal flow are possible, using an acoustically matched tube [1,2], but the tube interferes with lip and jaw movements. This interference is fatal to applications that require simultaneous recording of natural, undisturbed speech. One such application is the recording of a database for a text-to-speech-system (TTS). Most commercial TTS systems operate by piecing together segments of real speech. Natural, high quality recordings are a basic requirement, and one also needs to automatically and reliably estimate acoustic properties of the speech signals. Other estimates of glottal parameters, such as intubation to measure subglottal pressure [3], direct photography of the vocal folds [4,5], or photoglottography [6,7] to measure the glottal open area are also invasive and incompatible with sensitive applications. Plethysmography [8] can measure subglottal pressure, but is cumbersome, and the acoustic properties of the box in which the subject sits need to be carefully considered in order to get clean speech. Finally, Electroglottography (EGG) [9,10,11,12,13,14,15,16] is noninvasive, but measures vocal fold contact, rather than glottal flow; the folds can be contacting while the glottis is partly open, or non-contacting when the glottis is nearly closed. The EGG signal is only

weakly correlated with the width of the glottal opening or the glottal flow, and thus is not directly related to acoustic measurements.

Historically, the glottal flow has been estimated by “inverse filtering.” Inverse filtering is based on first estimating the vocal tract transfer function from a microphone or anemometer signal, inverting it, and using the inverted filter to remove the effect of the vocal tract [17]. If the estimated transfer function is accurate, the result will be close to the actual glottal flow.

However, the problem is intrinsically indeterminate, as one is trying to estimate both a time series of glottal flow and a time series of vocal tract parameters from a single input time series. In practice, assumptions are made that the microphone signal is quasi-stationary, that the spectrum of the glottal flow is simple (*e.g.*, a periodic pulse train, or a pulse train fed into a pre-emphasis filter), that the vocal tract can be modeled as an all-pole spectrum, and sometimes that the formant frequencies change smoothly with time. While the resulting algorithms work reasonably well, none of the approximations are perfect, and the resulting estimate of the glottal flow is only approximate, with ill-determined errors.

The goal of this study is to develop a new, noninvasive technique that allows an estimate of the volume flow of air through the glottis, U , based on an array of microphones and an EGG signal. We partially realize that goal by computing an estimate for U convolved with an unknown (but time-invariant) transfer function. In the remainder of the paper, this is what we mean when we speak of the QGG as an “estimator for U .”

This paper will first justify why such an algorithm is possible. Second, it will lay out the details of our algorithm. Third, it will describe a test of the algorithm. Since we do not have invasive flow measurements available for a direct comparison, we bring in two lines of indirect evidence to show that the Quasiglottogram (QGG) signal is closer to U than the standard signals used for voicing estimation. The first test is qualitative: we compare the QGG’s behavior to other signals in “difficult” regions of speech, and show that the QGG behaves well under conditions where one or another of the standard signals misbehaves. The second test is quantitative, though indirect: To prepare, we introduce a simple “toy” model of speech, and show that in that model, the signals that lead to the least variable amplitude estimates are the signals that are closest to U . We then show that the QGG signal allows a very steady estimation of amplitude, generally less variable than the result of other standard linear estimators. This provides evidence that the QGG signal is close to U .

A good estimator for glottal flow should be noninvasive. It should distinguish between voiced sounds and sounds generated in other constrictions of the vocal tract. It should also be linear, and should be related to the actual glottal airflow through a time-invariant transfer function. Linearity helps simplify connecting the estimate of U to physiologically important parameters like the subglottal pressure and the glottal open area, and it allows straightforward quantification of partial voicing. The output of the algorithm should be related to U via a time-invariant transfer function, so that we can meaningfully compare signals at one time to signals at another. Particularly, one would like to compare glottal flows between different phones.

We approach these goals by building a signal from linear combinations of several filtered microphone signals. If several microphones are placed near the head and neck, they will capture different signals, containing different information. We choose the linear combinations and filters to make the best possible match to the EGG signal, which provides an instantaneous

measurement of whether the vocal folds are contacting or not, and thus gives some indication whether the glottis is open or closed. We also use data from a microphone near the base of the throat to pick up a signal from the subglottal cavities. The subglottal cavities have an acoustic transfer function from the glottis through the throat wall to an outside microphone that is relatively independent of time, as (unlike the vocal tract) the trachea is not surrounded by muscles used to articulate speech. The dimensions of the trachea and bronchi are also largely unchanged during breathing, and even motions of the larynx are expected to lead to only modest changes in the acoustic resonances [18,19] of the trachea.

The algorithm requires at least two microphones, one to pick up the throat signal, and one for sound from the vocal tract. The throat signal contains a mixture of sound that propagated from the glottis down through the trachea then through the skin of the throat, mixed with sound that propagated from the glottis up through the vocal tract, then back down through free air to the throat microphone. The other microphones are used primarily to cancel out the component from the vocal tract. This cancellation results in a signal that can be thought of as the glottal waveform, filtered by its propagation through the neck. We show that this signal is less variable than commercial inverse-filtered signals, in the sense that the signal shows less phoneme-dependent variation. This stability is what one expects of the glottal waveform; it should be only weakly influenced by the vocal tract configuration [20].

It is necessary to cancel out the sound from the vocal tract because we want an estimator of the glottal flow, and we don't want our measurement to be disturbed by the dramatic changes in the transfer function of the vocal tract that occur during normal speech. This is a different approach from an inverse-filter estimator, which attempts to dynamically estimate and invert the vocal tract transfer function: a nontrivial, multiplicative operation on one signal. We look for a signal that has a time-invariant relationship to the glottal flow via a linear operation on several signals.

2. Experimental Method

In the experiments described below, we mounted four Brüel and Kjær type 4165 omnidirectional microphones on the face guard of a hockey helmet. The microphones were mounted near the nose (4 cm lateral from the end of the nose), to the side of the mouth (2 cm lateral of the corner of the mouth), near the forehead (centered, 11 cm above nose level), and near the base of the throat (on centerline, 2 cm from skin, 2 cm above the top of the sternum). The microphones were checked to be slightly outside breath streams, and were protected by 4 mm of windscreens. In other experiments, we found that only the placement of the throat microphone was critical: it should be placed as close as possible to where the trachea can be palpated (the *fossa jugularis*), so long as the microphone does not collide with the subject during normal head motions. The placement of the other microphones is not critical, and we have obtained similar signals from a 6-microphone array (including cheek and back-of-the-head microphones), and a 3-microphone array (using a gradient microphone for the throat signal).

The EGG signals were obtained from a Portable Laryngograph [21]. The electrodes were placed to maximize the signal strength for long vowels, in modal speech, over the normal range of f_0 . Data was digitized at 12 kHz per channel with an antialiasing filter. The microphone and EGG signals are high-pass filtered at 40 Hz to reduce room noise and the large amplitude, slow

components of the EGG that correspond to motions of the larynx. All inverse filtered signals were produced by ESPS/waves [22]. They are derived from microphone **M** (near the mouth). ESPS/waves computes the linear predictor all-poles filter that will whiten each frame of **M**, then inverts the filter and applies it to **M**.

In the examples that follow, we used Mandarin speech acquired from a female subject, a native Mandarin speaker fluent in English. Because Mandarin is a tone language, Mandarin speakers can conveniently control f_0 movements. Therefore we expect that vocal fold vibration will be more repeatable, reducing unpredictable fluctuations in U .

The data presented in section 4 was taken from a random (*e.g.*, selected for other purposes) set of English vowels and words, read at normal volume in modal speech, along with some unplanned spontaneous phrases. The data for section 5 is a database of 979 utterances in the form “*Ta shuo X san tian.*” (meaning “He says X for three days.”). In the database, the frame syllables all have high level tone (tone 1) and X is a randomly selected 979 of the ~1300 allowable Mandarin syllables (here we include the tone as part of the syllable). The selection was made without duplication. The subject pushed a key for each prompt, and Chinese characters were presented on a screen. Speech began 0.72 ± 0.1 seconds after the prompt. Overall, f_0 had a mean of 235 Hz and a standard deviation of 26 Hz, with most of the variation on X .

Calibration utterances to train the algorithm (*i.e.*, fit the filter coefficients), for both sections, were a broad mixture of sounds, totaling 197.3 and 232.5 seconds for the two sessions. Half were single-syllable English words (*e.g.*, “nap”, “sap”, “bee”, “gold”, “hold”, “moo”, ...). The words provided a variety of vowels and consonants in natural contexts. One quarter of the utterances were repeated, unvoiced fricatives (“ch”, “f”, “s”, “sh”, “th”) with no vowel context, and repeated, aspirated “p”, and “h” (again, no vowel context). These voiceless sounds are useful to force cancellation of the mouth signal and to check the operation of the algorithm, because their U is a relatively steady flow, compared to voiced speech. We used 25 calibration utterances from the beginning of the day’s recording (randomly selected from ~40 recorded utterances), and another 25 at the end, ~6 hours later.

The remainder of the calibration utterances were long vowels, nasals, and voiced fricatives (“a”, “u”, “v”, “z”, “m”) where f_0 was swept either up or down. The speaker was asked to sweep at least half of her pitch range in each utterance, catching either the lowest or highest comfortable pitch. These f_0 sweep utterances provide voiced speech with the speaker’s full pitch range. It is useful for the calibration utterances to cover all the f_0 range of the speech used in the experiments. The calibration utterances did, having mean f_0 of 248 Hz and standard deviation 97 Hz. The sweep rates for the sweep utterances were 150 ± 70 Hz/s (average slope over the central 1 second of the utterance).

3. The algorithm

The overall structure of the algorithm is shown in Figure 1. To explain the design of the algorithm, we can consider a simplified form of the microphone array as shown in Figure 2: one microphone, **M**, near the mouth, and another, **T**, at the base of the throat. During voiced sounds, the signals are excited by the flow, U , through the glottis. The signal at **T** is made of two main

components, one traveling through the vocal tract (transfer function V) to \mathbf{M} , the other traveling directly through the neck [23] (*via* a transfer function N). Note that N includes the resonances of the trachea, in addition to the transmission through the throat (which acts as a low-pass filter). Sound coming out of the mouth, and then propagating down through the air to \mathbf{T} has a transfer function $A \cdot V$.

In this paper, we treat these transfer functions as generic matrix manipulations in any complete basis, not yet specializing to a time series, frequency representation, or some intermediate basis, following [24]. The throat microphone signal is then $T = (N + A \cdot V) \cdot U$. We take $M = V \cdot U$, neglecting the small amount of sound that comes out through the neck and propagates upwards through the air to \mathbf{M} .

We expect that N and A should be nearly independent of the phoneme being produced. On the other hand, V varies dramatically and systematically as a function of the phoneme. We can express this variation by writing $V = \bar{v} + \alpha \cdot \tilde{v}$, where V depends on the phoneme only through α , and $\langle \alpha \rangle = 0$ (the angle brackets denote an average over the corpus of speech). Because human speech is the result of several independently controlled articulators, α is normally a vector. In other words, we decompose V into a constant transfer function, \bar{v} , and a variable part. Then, $T = (N + A \cdot \bar{v}) \cdot U + \alpha \cdot A \cdot \tilde{v} \cdot U$.

Hypothetically, if we had a reference signal, Q , which was derived from U via some time-invariant transfer function, q , we could find a linear filter that would reproduce Q from the microphone signals. (Note that we assume off-line processing, so our filters need not be causal.) The general form for that linear filter would be $C = c_M \cdot M + c_T \cdot T$, and to find the filter matrices (coefficients) c_M and c_T , we would minimize the difference between C and Q :
 $c_M, c_T = \arg \min E(c_M, c_T)$, where

$$\text{Eq. 1} \quad E(c_M, c_T) = (C - Q)^T \cdot (C - Q) = \|C - Q\|^2$$

is the sum of squared errors between our target signal Q and our reconstruction.

If $\alpha = 0$, Eq. 1 is degenerate, and its solutions are all filter coefficients c_M and c_T on a line which we will call \mathbf{L} , that goes through $c_{M1} = 0$, $c_{T1} = q \cdot (N + A \cdot \bar{v})^{-1}$ and $c_{M2} = q \cdot \bar{v}^{-1}$, $c_{T2} = 0$. The first of these two solutions corresponds to using a single filter on the throat microphone to match Q , while ignoring M . The second solution is the reverse: using a single filter on the mouth microphone while ignoring N . Points on \mathbf{L} correspond to linear combinations of these two filters, and all points on \mathbf{L} are equally good solutions and match $C = Q$ exactly, so long as $\alpha = 0$, assuming that the necessary inverses exist. This derivation can be extended to allow white additive noise in the microphones, in which case the necessary inverses always exist. The results are qualitatively similar, although the derivation and results become substantially more complex.

If the speaker starts talking, instead of just vocalizing with a stable vocal tract, the transfer functions will vary from phoneme to phoneme, and we will not be able to match Q perfectly at

every moment, so E will be positive. Not every solution gets the same increment of error, though. Solutions that depend predominantly on V will fit worse and have larger errors than solutions that depend predominantly on N , because V changes from phone to phone. This difference breaks the degeneracy and typically picks out a single best solution. In general, the best solution will use signals from all microphones, and it will provide a better approximation to Q than could be obtained from any linear filter operating on any single microphone in the array.

If we assume that V varies slowly enough, we can write down the change in error due to the difference between V at a given moment and the average of V (*i.e.*, \bar{v}):

$$\text{Eq. 2} \quad \Delta E = \left\langle \left\| (c_T \cdot A + c_M) \cdot \alpha \cdot \tilde{v} \cdot U \right\|^2 \right\rangle,$$

where the average (written as angle brackets) is taken over all phonemes in the corpus. The change in error is always non-negative and is normally nonzero everywhere except on a hyperplane we will call \mathbf{P} , defined by $c_M = -c_T \cdot A$. Not coincidentally, this relationship between the filter coefficients is exactly what is needed to cancel out the part of the signal that came from the mouth, leaving only the part that came through the throat wall. The intersection of \mathbf{L} and \mathbf{P} then specifies the best estimator for Q (the one that is least sensitive to changes in the transfer functions). These results generalize to arrays of more than two microphones. They also can be generalized to the case where all the transfer functions vary with time, though Eq. 1 and Eq. 2 will change in detail.

So, given a reference signal, we can start with an array of microphones, and find the linear combination of filtered microphone signals that best matches the reference signal. The resulting signal is less variable than any signal derived by a linear operator from a single microphone. Loosely speaking, the best estimator is obtained by canceling out the highly variable signal from the mouth microphone, and using the part of the signal that did not propagate through the vocal tract.

In the real world, one doesn't have a perfect reference signal, Q . The best we can obtain non-invasively is the EGG signal. The EGG is related to U in a nonlinear and variable manner, because the larynx moves up and down relative to the electrodes used to measure the EGG. Repeating the above analysis shows that the variability of EGG measurements are not important, so long as changes in the EGG signal are not correlated with phonemes. This is true, by and large, as the larynx moves in response to pitch changes and inhalations, neither of which are correlated with most phonemes in most languages. Glottal stops and pharyngeal sounds are an exception, but they are not particularly common, typically comprising just a few phonemes in a language.

Nonlinearities in the relationship between U and the EGG signal are difficult to analyze analytically. We have investigated their effect empirically, in Sections 4 and 5.

3.1. Algorithm Introduction

We first build the data matrix, X , where each row contains the signal from one of the microphones, and each column corresponds to one moment in time. X is an n by m matrix when

there are n microphones, each digitized to produce m samples of audio. The EGG signal is a one by m matrix.

We then select a set of taps (*i.e.*, taps on a delay line) for each microphone. Physically, the closure of the glottis is the cause of the acoustic signals: when it closes, a sound wave propagates up the vocal tract and down the trachea, reaching the microphones roughly a millisecond later. If one tries to observe a glottal closure at time t , one will need to use microphone data from later times, when the sound waves from the closure reach the microphones. Thus, we use the taps to build a finite impulse response (FIR) filter to predict the present EGG signal from future microphone signals.

To select the span of the taps, we need to consider the purposes of the filter we are building. It needs to cancel out the mouth signal that is picked up by the throat microphone, and it needs to match the impulse response of the remainder of the throat microphone signal to the EGG signal.

The taps span the range of delays beginning with the earliest propagation from the glottis to the microphone in question, ending when the impulse response of the vocal tract goes below 1% of its peak value.

Vocal tract formant bandwidths can be as small as 40 Hz [25,26] when the glottis is closed, though in real speech a bandwidth of 80 Hz is more realistic [27]. Such a bandwidth implies that the vocal tract resonances will take about $3 \cdot 1/(2\pi B) \approx 6\text{ms}$ to decay. Bandwidths for the tracheal resonances are wider, 200-400 Hz [28], and so are not the limiting factor for the window width in the time domain. Matching the acoustic to the EGG signal requires a window length of about $1/f_0$, which can be slightly longer. We choose the range of taps to cover the longer of these times (13 ms in these experiments), and we use the same number of taps for every microphone.

One needs enough data to capture a wide range of speech conditions, especially speech with a range of fundamental frequencies (f_0). The algorithm will not produce good estimates for speech that has f_0 outside the range of the training data in x .

The QGG signal does have some imperfections. It has a nonzero (though relatively small) response to fricatives and plosives. Also, the QGG algorithm does not estimate U directly, but estimates U times a transfer function, where the transfer function depends on the subject and the experimental configuration. Finally, the algorithm does not adapt to changes in microphone positions, relative to the mouth, nose, and throat: stably mounted microphones are presumably required for good results.

Note that we do not claim any absolute calibration for the QGG signal. Because the QGG filter coefficients involve the EGG signal, the QGG amplitude will differ from person to person and session to session, depending on neck structure and the placement of the EGG electrodes. However, the QGG signal depends on the EGG only through its average properties during the calibration/training session. So, when one is actually using the QGG (as opposed to calibrating it), the EGG is entirely unused, and may be disconnected. Therefore, factors that affect the EGG signal during use (such as the momentary position of the larynx with respect to electrodes) will have no effect on the QGG.

There are similarities between the QGG signal and the low-pass filtered mouth signal. Both are generated by linear, time-independent filters. Indeed, in practice, the QGG signal also has less high-frequency power than microphone **M** (though more than the output of a high-order low-pass filter). The steep QGG spectrum has its origins in three places: the EGG signal that it is matching has relatively less high-frequency power than speech, and high frequency components radiated from the mouth tend to have phases that are strongly dependent on the phone, consequently there is little cross-correlation between the high-frequency EGG components and the high frequency components from **M** and they are de-emphasized. Finally, microphone **T** picks up some signal with a stable phase relationship to the EGG, but high frequencies are not effectively coupled out through the throat wall.

However, we expect that the QGG should give a more consistent amplitude measurement than the low-pass filter. Since the low-pass filter signal essentially measures the amplitude of f_0 in microphone **M**, it will be strongly boosted when the first formant is close to f_0 . The QGG signal should not be affected by this problem, as it is measuring a combination of signals chosen to match the EGG's relatively stable amplitude: roughly speaking, it is measuring the subglottal cavities through the throat wall, and the formant structure of the vocal tract has little effect down below. The QGG ought also to be more consistent than a low-pass filter if the speaker has a pitch range that swings through more than an octave. In that case, the low-pass filter will sometimes pass the second harmonic, and sometimes not, leading to more amplitude fluctuation.

3.2. Implementation

We build a set of linear equations corresponding to the FIR filter that best predicts the EGG signal, using straightforward least-squares linear prediction techniques²⁹. The predicted EGG signal at each time is a linear combination of $n \cdot q$ values (q taps on each of n microphones). To start, we define the covariances between shifted signals from the i^{th} and j^{th} microphones:

$$\text{Eq. 3} \quad \phi_{\alpha,i,j} = \frac{1}{m} \sum_t x_{i,t-\alpha} \cdot x_{j,t},$$

where α is the time shift between the i^{th} and j^{th} microphones (we neglect end effects, for simplicity), and t indexes the time. These covariances are estimated from the data, and contain noise covariances. Analogously, we will write $\phi_{*,\alpha,j}$ for the covariances between the EGG signal and the shifted microphone signals. The filter coefficients that minimize the mean squared error are then the solution to

$$\text{Eq. 4} \quad \sum_i \phi_{\alpha,i,j} \cdot c_{\alpha,i} = \phi_{*,\alpha,j},$$

which is a set of $n \cdot q$ linear equations. We prepare to solve the equations by stacking the $c_{\alpha,i}$ to make a single vector C , stacking $\phi_{*,\alpha,j}$ to make a single vector P (of size $n \cdot q$ by 1), and placing the elements $\phi_{\alpha,i,j}$ into the corresponding places $H_{\alpha+i,j}$ to make a $n \cdot q$ by n matrix. We solve the resulting matrix equation, $HC = P$, with a singular value decomposition algorithm

to allow for degeneracies and near-degeneracies. C can then be unpacked to yield the $c_{\alpha,i}$, which are the filter coefficients that give the best prediction of the EGG signal from the set of microphone signals.

We can now calculate the QGG signal,

$$\text{Eq. 5} \quad p_t^{[1]} = \sum_{\alpha,i} x_{i,t-\alpha} \cdot c_{i,\alpha}$$

from the $c_{\alpha,i}$ and microphone signals. The prediction is not at all precise because the EGG signal is a strongly nonlinear function of the glottal opening: it contains little information beyond the simple fact of whether the vocal folds are touching or not. It would be surprising indeed if one could build a linear filter that would exactly match the EGG signal.

3.3. Examples

Figure 3 shows the raw data for all four microphones and the EGG signal for a small section of speech. The display is centered about the burst of a “t”, and shows voiced regions (at the edges), a burst, and then a fricative-like region after the burst. Note that the voiced regions are quite similar in all four microphones (though the throat microphone, **T**, is phase shifted by about 30°). On the other hand, the fricative-like region (3.44 s – 3.46 s) has a substantially different structure at microphone **T**, and also has an amplitude that is relatively small, about 3 dB lower, relative to the voiced regions than the other three microphones.

Figure 4 shows the output of the QGG and several comparison signals in the same region around the “t” burst. Figure 5 shows the QGG and comparison signals on a longer stretch of speech.

4. QGG as fodder for voicing estimation.

An engineering evaluation of the QGG signal as input for a voicing estimation algorithm is beyond the scope of the paper. Instead, we will show that (at least under some conditions) the QGG signal can provide a more reasonable indication of the presence of voicing than either the EGG or the inverse-filtered mouth signal. This is the first, qualitative, test of the QGG.

The advantages of the QGG signal follow from its construction: because it is a linear function of the pressure near the glottis, it is well behaved during startup and shutdown of the glottal oscillator. So, unlike the EGG, it may be able to quantify partial voicing, and mark onsets of voicing precisely. Because the QGG is constructed from a time-invariant filter operating on acoustic signals, it may be more robust than algorithms based on an inverse filter (we do not discuss manual adjustment of inverse filter coefficients here, as such techniques are impractical for large speech corpora). Any time the spectral estimation step of an inverse filter fails to produce a good result, or any time the speech signal is not well represented by an all-poles

transfer function, one expects the inverse-filtered signal will not reflect the glottal state. The QGG doesn't suffer from those problems.

Figure 6 and Figure 8 show examples of speech signals where the glottal oscillation is starting or stopping. The figures show that the QGG signal can sometimes provide a much better explanation of the acoustic signal than does the EGG signal. Limitations of EGG signals have previously been described elsewhere [30,31].

In Figure 6, the envelope of the QGG signal tracks acoustic power (the mouth signal), while the EGG signal shows an unnaturally sharp onset/ending. If the glottal oscillation stopped with the EGG signal, one would have to assume a bandwidth for the first formant of only 5 Hz for the acoustic signal to persist as long as it does, which is incompatible with the known bandwidths [25,26,27]. The vocal tract simply isn't a good enough resonator for the sound to persist 30 ms after the end of glottal oscillation. Therefore, glottal oscillation must be continuing at a lower level (so the vocal folds don't touch) without showing up on the EGG. Titze [32] and Stevens [33] have discussed this kind of small oscillation.

Other observations can also be explained similarly, such as the events around $t \approx 8.015$ s in Figure 6. Imagine comparing two similar vowels, one uttered with amplitude just small enough so that the glottal folds don't collide, and the other uttered with slightly more amplitude so that the folds do collide on each cycle. If we consider a decomposition of the signal into a stack of harmonics at $f_0, 2f_0, 3f_0, 4f_0...$ the lowest harmonic will primarily measure the total air flow per cycle, and will change only gently and continuously when the vocal folds begin to collide.

However, the higher harmonics do not behave smoothly. Below the collision threshold, the oscillation is close to a simple harmonic oscillator, and there is little power in the harmonics. Above threshold, there is a strong nonlinearity when the vocal folds collide; large amounts of power suddenly start to be generated in the harmonics. This is the situation that is described by the two-mass model [34,35], which typically gives a spectrum where the amplitude of the n^{th} harmonic scales as $a_n \propto n^{-1}$, or a 6db/octave slope. Such a change can be seen in the spectra shown in Figure 7.

The effect is not confined to the low-amplitude tails of voiced sounds. For example, in Figure 8, an acoustic signal begins two periods before the first EGG activity. Again, one must assume that the vocal folds are oscillating but not yet completely closing.

Figure 9 shows a section of a low-amplitude, sustained 'o' as an extreme example where the EGG fails to explain the acoustic signal. Several times, the amplitude of the EGG signal jumps up dramatically for a few milliseconds. Little effect is seen in the acoustic signal, other than an increase in the power of the higher harmonics, nor is much change seen in the QGG signal. Hypothetically, the vocal folds could be oscillating within a whisker of touching, and some perturbation intermittently reduces the spacing just enough to make them collide. The noteworthy observation here is that the QGG signal is a better predictor than the EGG of whether the acoustic signal shows voicing. This should not be too surprising, as the QGG is constructed from a set of acoustic signals.

These problems we have displayed are not hard to find, occurring at these levels in 6 of 304 voiced syllables inspected. Because the problems seem to be associated most with glottal oscillations where the vocal folds do not contact, we expect that EGG problems should be much more common in languages that make more extensive use of murmurs (*i.e.*, a 'breathy' voice quality), most notably Hindi. The inverse filtered signal also tends to behave badly for low amplitude voicing or other conditions where the power in the higher harmonics is very low. Among the displayed signals, the QGG provides signals that display a strong contrast between voiced and unvoiced regions, and have most of their power in the fundamental to reduce the likelihood of octave errors in any pitch tracker to which the QGG signal is fed.

As input to a voicing estimator, many of the advantages of the QGG signal are shared with the low-pass filtered mouth signal. However, we would expect that the suppress bursts (*e.g.*, p, t, k) and be a closer representation of the state of the larynx during closures.

5. QGG as a measure of amplitude/emphasis

The Quasi-glottogram signal is valuable for more than correcting voicing errors. It also provides an estimate of the amplitude of the oscillatory flow through the glottis. We expect that this amplitude will be a better predictor of prosodic emphasis and a better measure of the speech effort being expended by the speaker than is the total acoustic power of the mouth signal, the inverse filtered mouth signal or the EGG.

Amplitude has been known to be a significant component of prosody since the 1950s [36,37,38,39,40], and into more recent literature [41,42,43,44,45]. However, all these studies have been severely limited by the large intrinsic variability of speech amplitude measurements. The experimental designs (*e.g.*, ANOVA analysis on p. 190 of [41]) invariably compare the prosodic effect in question to the unpredictable variations. Reducing this variability can be seen to be just as good as having a larger effect to measure. This is one value of the QGG: it allows a cleaner, low variance amplitude measurement, and should lead to more conclusive experiments.

5.1. Model of amplitude variance

To justify our intuition that the QGG signal will allow better amplitude measurements, consider a toy model of the speech apparatus: a glottal source that drives the vocal tract, which we model as a time-varying filter. Loosely speaking, the variability of the amplitude outside the mouth comes from two sources: intrinsic variability in U , and changes in the coupling through the vocal tract transfer function, V . Since the two variances add, the variability of the mouth power will be greater than the variability of glottal power. Consequently, we expect that the best linear estimators of the glottal source should have the lowest variability. We can use this as a figure of merit to compare algorithms: less variable estimators are better and probably closer to the glottal signal.

We will work in a short-time Fourier transform basis to conveniently describe speech-like signals. Signals are then indexed with two parameters: a time index, t , which locates the transform's window, and a frequency index, ω , for the low-resolution spectrum in the window. In this toy model, the glottal source changes its amplitude but not its spectral shape:

$U(t, \omega) = f(t) \cdot g(\omega)$, where we assume that the amplitude, $f(t)$, changes slowly compared to structure in the glottal spectrum, $g(\omega)$. We can safely assume $\sum_{\omega} |g(\omega)|^2 = 1$ without affecting the model, as the overall amplitude can go into $f(t)$.

Next, we can write a time- and frequency- dependent transfer function for the vocal tract: $h(\omega, \phi(t))$, where $\phi(t)$ is the vocal tract configuration (roughly, the phone) at time t . The pressure outside the mouth is then $s(t, \omega) \approx f(t) \cdot g(\omega) \cdot h(\omega, \phi(t))$, and we can sum over frequency to get the RMS amplitude of the mouth signal: $a^2(t) = \sum_{\omega} |s(t, \omega)|^2 = f^2(t) \cdot z^2(\phi(t))$,

where $z^2(\phi) = \sum_{\omega} |g(\omega) \cdot h(\omega, \phi)|^2$ shows how efficiently power is coupled from the glottal source out through the mouth for a particular phone ϕ .

We can now take the log of the power to write

Eq. 6 $\log(a(t)) = \log(f(t)) + \log(z(t)).$

As long as variations of the vocal tract are uncorrelated with changes in the larynx, the variances of the two right hand terms add, and we can conclude that $\text{var}(\log(a)) = \text{var}(\log(f)) + \text{var}(\log(z))$. Since $\text{var}(\log(z)) > 0$, $\text{var}(\log(a)) > \text{var}(\log(f))$. In other words, the amplitude outside the mouth is always (in this toy model) more variable than the amplitude of U . The same conclusion follows if you consider h to include the vocal tract plus an arbitrary linear operator: thus any filtered version of the mouth signal will still be more variable than U .

The limits to the assumption that f and z are uncorrelated come from two sources: First, the speaker's intentional muscle motions can have correlations between the glottis and the vocal tract (e.g., a hypothetical language might specify that high vowels are always spoken in a pressed voice). Second, some vocal tract configurations with tight constrictions can change the glottal waveform. However, neither circumstance seems common.

As a concrete example of this, consider a vowel where f_0 matches the first formant frequency, F_1 . Acoustic power will then be efficiently coupled from the glottis out the mouth, because a peak of $g(\omega)$ matches with a peak of $h(\omega, \phi(t))$, z will be large, and the amplitude at the mouth will therefore be large. On the other hand, if $\frac{3}{2} f_0 = F_1$ with the same amplitude of U , the fundamental frequency will be below the first resonance of the vocal tract, z will be small, and the amplitude at the mouth will be small. Comparing the two cases, we see that substantial variance in amplitude can be generated as the sound wave propagates through different configurations of the vocal tract.

This toy model contains several loose assumptions and shouldn't be taken too far, but it does give important clues for finding good algorithms, since the mathematics remains valid if the

transfer function, h includes the behavior of the microphones and a linear signal processing algorithm. For example if h is time-invariant, $\text{var}(\log(z))$ will be zero, and the variability of the amplitude estimate, $\text{var}(\log(a))$, will be as small as possible. Conversely, if s is always close to g , the transfer function must be near unity, so that $\text{var}(\log(z))$ must be small and thus $\text{var}(\log(a))$ will be as small as possible. So, we expect that algorithms that are good glottal estimators will give stable amplitude estimates and vice versa. Note that for further arguments, we do not require any of the details of the toy model, merely this conclusion, which is independently testable, and likely to be true independent of the model.

Purely pragmatic considerations will also lead us to the same figure of merit. If one is studying prosody, then any variation of amplitude that is a function of the phoneme should be considered noise: it prevents one from comparing prosodies of different words. Good comparisons are only possible if the amplitude measure is independent of phoneme. So, we would like to improve the signal-to-noise ratio of prosody experiments by reducing the noise, which again means finding an amplitude estimator that is less variable.

5.2. Comparison of QGG & acoustic power variance

We conducted a second, quantitative test. This test directly establishes the usefulness of the QGG signal as a measure of amplitude prosody. We also show that the QGG signal allows a very steady estimation of amplitude, less variable than the result of other standard linear estimators. Following the logic in Section 5.1, this test indirectly establishes that the QGG signal is a reasonable estimator of a filtered version of U .

We used the database of 979 utterances in the form “*Ta shuo X san tian.*”, described above. We calculated the QGG for all the utterances, along with an inverse-filtered \mathbf{M} , low-pass filtered \mathbf{M} , and the raw \mathbf{M} . The boundaries of the variable syllable (X) were hand-segmented, and an algorithm (ESPS/waves *get_f0*) was run to find the two voiced regions on the left side of the segmented area (assumed to correspond to the vowels in “*ta*” and “*shuo*”), the voiced region inside X , and the two voiced regions to the right of X (assumed to correspond to the vowels in “*san*” and “*tian*”). Four of the utterances were voiced through between “*ta*” and “*shuo*” and were dropped, because we could not assume that the voiced region matched the vowels. We then calculated the mean power near the center of the five voiced regions, using a cosine window.

Table I shows the standard deviation of $\log(\text{power})$ for each combination of region and signal. In every case, the QGG signal is more reproducible than the others, yielding (on average) a standard deviation 46% smaller than the corresponding low-pass filtered speech, 38% smaller than unprocessed speech, and 20% less than the inverse filtered signal. These improvements in SNR are conservative limits, as the speech contains some intrinsic variability that cannot be removed by signal processing.

Comparisons of cells in Table I have several implications. First, one expects the variation of the frame (regions 1, 2, 3 and 4) to be small, where we always have the same syllable in the same position. In contrast, the variation in the X region should be bigger, because the syllable identity changes. This is reflected in Table I, where the X region shows the largest variation under all conditions.

Why is the QGG variability not smaller in the X region, if it indeed removes the effect of changes in the vocal tract? It is not smaller because, unfortunately, the QGG is not an estimator for U , but rather (see section 3) a filtered version of U , $q \cdot U$, where the filter, q , is time-invariant but can vary as a function of frequency. The QGG estimate can therefore vary with the f_0 of the speech.

We tested this by choosing a subset of the data that all has similar f_0 , and re-calculating the standard deviation of region X. We chose all 260 syllables that have Mandarin tone 1. This is a high, level tone, which is the same tone that occurs in regions 1, 2, 3, and 4. In this subset of the data, the f_0 trajectory is relatively flat in each syllable, as well as across the whole utterance. The results are displayed in column X' of Table I. The standard deviation of the amplitude of all the signals is lower, but the QGG drops most dramatically.

One can also see both the pitch-dependence of the QGG and its relative insensitivity to the vocal tract configuration in Figure 10. This is a scatter plot of the mean f_0 of all syllables vs. the measured amplitude of the QGG signal. We used the same window (as defined above) to calculate the mean f_0 as was used for the amplitude measurement. Syllables with tone 1 are seen as a tight cluster in the upper right corner. That cluster spans the full range of phonemic variation, covering all vowels in combination with a variety of consonants and glides. The other syllables in Mandarin (shown as dots) have either low pitch, or they are rising from or falling to low pitch, thus they have an average pitch below that of tone 1. There is a clear trend of increasing amplitude measurement with increasing pitch, perhaps as the result of q . One could empirically correct for this trend, if one knew how much of the effect is the result of q , and how much is the result of the speaker's glottal flow changing as a function of pitch.

While the QGG algorithm dramatically reduced the variation of Regions 1 and 2 (to around 10% in table I), it was less successful in Regions 3 and 4. This suggests that there is more inherent variability in Regions 3 and 4, which may well be some carry-over from the pitch and phoneme change in region X. This larger inherent variation can also be seen for all the signals, although not as clearly because the other algorithms don't yield as reproducible an amplitude signal as the QGG.

It is possible that much of the variability in Regions 1 and 2 is external to the speaker. The major source of variability outside the speaker is expected to be the motion of the throat microphone relative to the throat if the speaker tilts the head forward or back, because the throat microphone was fixed to the helmet. We do not have accurate estimates of this effect, as head motions were not measured. However, we estimate (after the fact) that 2 mm changes in throat to throat-mic distance are quite plausible, and these could yield amplitude changes of 5%-10%. A secondary effect is that the design of our microphone array allows the microphones to move about 1 mm relative to the skull. This movement is expected to lead to ~3% changes in values of the transfer functions (changing some more than others), which would lead to comparable changes in measured amplitude.

5.3. QGG for amplitude prosody

As a further test with the same database (Table II), we attempted to eliminate any changes in amplitude by predicting the amplitude of each region in terms of the amplitudes of the other regions. Essentially, this normalizes the measured amplitudes to the rest of the utterance, and would eliminate the effect of a uniform change in amplitude from utterance to utterance. The goal here is to reduce the variability deriving from the experimental subject (*e.g.*, from changes in the volume of inspired air) and focus more tightly on variations that result from the signal processing.

We fit a least-squares linear predictor to the logs of the amplitudes, *e.g.*,

$$\text{Eq. 7} \quad \log(\hat{A}_X) = b_0 + b_1 \cdot \log(A_1) + b_2 \cdot \log(A_2) + \dots,$$

and then measure the RMS size of the residual, $\log(A_X) - \log(\hat{A}_X)$. Here, A_R is the amplitude of the R^{th} region (a RMS average of the signal inside a cosine window), and \hat{A} is the predicted amplitude, based on the other regions. The predictor is a five-parameter linear fit, and we fit separate predictors for regions 1, 2, X', 3, and 4. Again, the QGG leads to a nicely small variance: it has predictable amplitude. Surprisingly enough, the low-pass filtered signal is comparably predictable, even though its performance before prediction (Table I) is quite poor.

The most important terms are generally those that predict A_3 in terms of A_4 and vice-versa. After prediction, the frame regions on both sides of X have similar variabilities, approximately 8% for the QGG signal. This remainder seems intrinsic to the speaker. The excess variation in regions 3 and 4 is gone, even though the variation in regions 1 and 2 is practically unchanged. This suggests that amplitude variations in X, which are driven by phoneme and f_0 changes, carry forward into the following syllables.

Because region X' contains a diverse set of different syllables while the frame regions (regions 1, 2, 3, 4) always have the same syllable in the same position, we expect more variability in the amplitude of X'. Table II shows this. However, changing syllables only disturbs QGG amplitude measurements by 11% (beyond the intrinsic 8%), so long as the pitch is reasonably stable.

6. Summary

We have shown that the QGG algorithm can produce a useful, noninvasive estimate of the glottal flow. (More precisely, the glottal flow filtered by an unknown linear filter.) When used to analyze speech, it can be well behaved under conditions where EGG and inverse-filtered signals would lead to errors in voicing estimation. It also yields substantially more stable amplitude measurements than other techniques. Amplitude measurements of repeated words using the QGG signal can have a fractional standard deviation as small as 9%, and the fractional standard

deviation across different syllables in the same location in the sentence is just 37%, with much of that standard deviation being attributable to intonation differences.

The QGG algorithm should find applications in studies of the amplitude part of prosody. We also see applications in text-to-speech systems, where there is a need for reliable automatic processing of speech data, and possibly in medical screening or diagnostics of voice disorders.

We acknowledge Robert Kubli and Gary Elko for equipment and advice, Mohan Sondhi and Joseph Hall for discussions and Ronald J. Baken and three reviewers for substantial comments.

Tables

Fractional standard deviation of amplitude	Region 1 <i>ta</i>	Region 2 <i>shuo</i>	X <i>X</i>	X' (tone 1 only)	Region 3 <i>san</i>	Region 4 <i>tian</i>
Raw	0.22	0.16	0.48	0.32	0.31	0.29
Low-pass Filtered	0.29	0.23	0.42	0.34	0.32	0.35
Inverse Filtered	0.14	0.14	0.42	0.24	0.22	0.24
QGG	0.09	0.09	0.37	0.16	0.19	0.20

Table I: Variability of speech amplitude, after processing by four algorithms. Region X is the variable syllable, and shows larger variability because the amplitude is a function of the phoneme in X.

Fractional standard deviation of amplitude, after linear prediction.	Regions 1 & 2 <i>ta shuo</i>	X' <i>(tone 1 only)</i>	Regions 3 & 4 <i>san tian</i>
Raw	0.16	0.32	0.11
Low-pass Filtered	0.09	0.16	0.06
Inverse Filtered	0.14	0.23	0.10
QGG	0.08	0.14	0.09

Table II: Unpredictable variability of speech amplitude, after processing by four algorithms. Here, we use the subset of the data where X has tone 1 (a high level tone), so that the pitch matches regions 1, 2, 3, 4. The amplitude in each region was predicted in terms of the other four regions; the table shows the remainders.

Figures

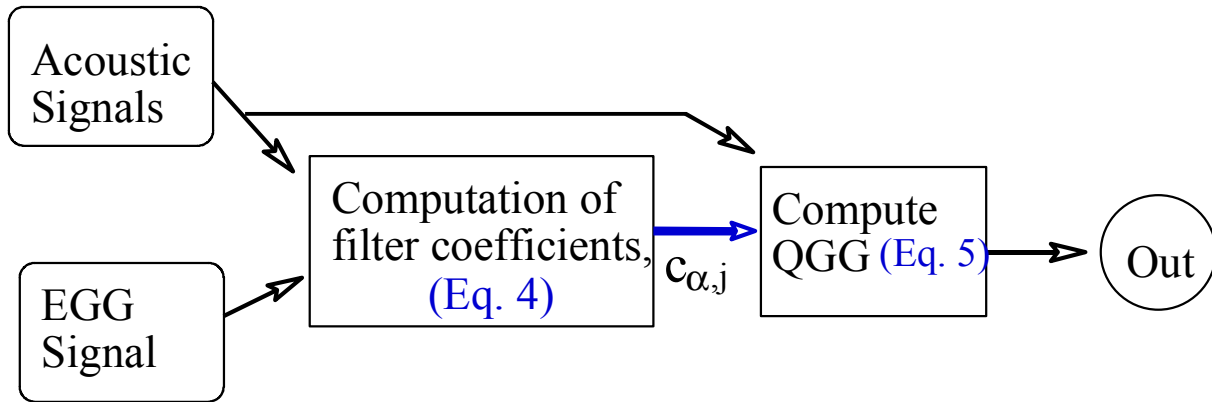


Figure 1. Diagram of the data and signal flow in the computation of the QGG.

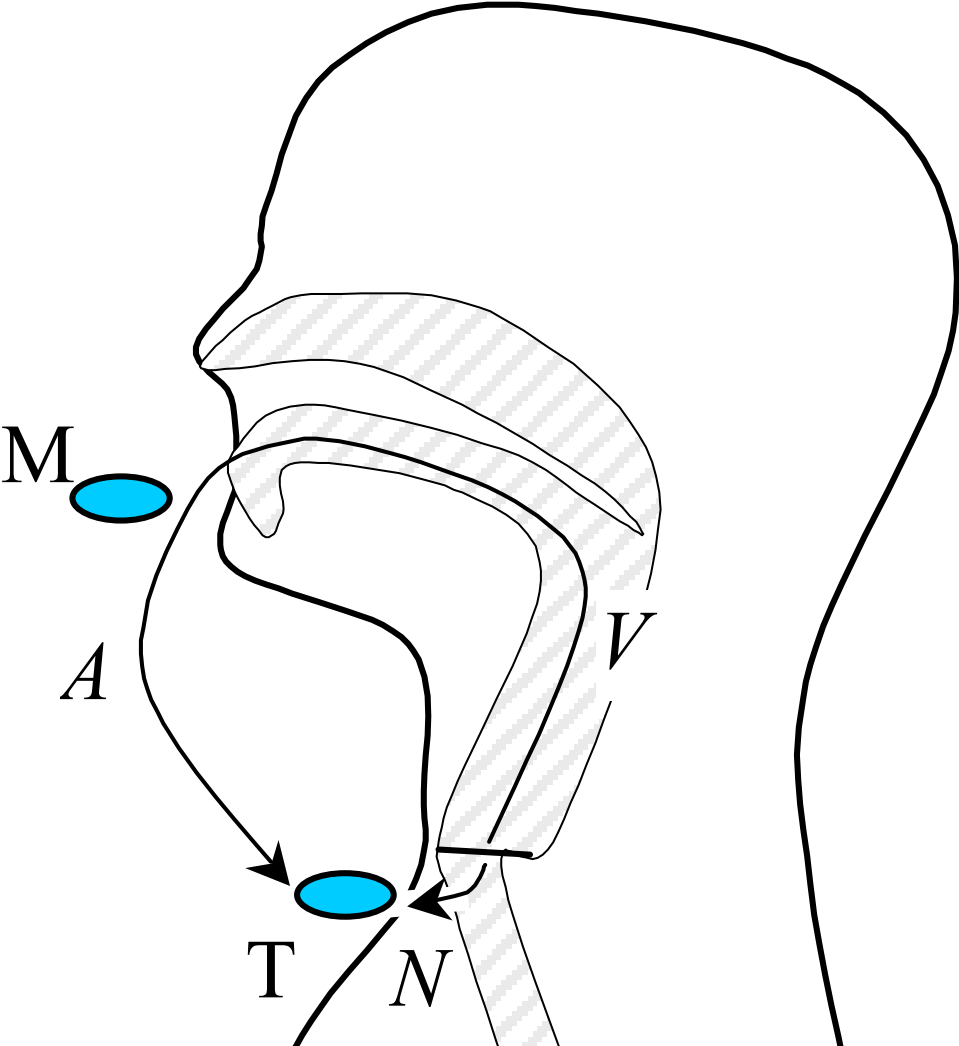


Figure 2. Schematic of microphone placement and signal paths. **M** is a microphone near the mouth; **T** is a microphone near the throat. V refers to the transfer functions from the larynx along the vocal tract to **M**, and $A \cdot V$ is the transfer function from the larynx, through the vocal tract, to **T**. N is the transfer function from the larynx through the throat wall to **T**.




Figure 3. Raw signals, used as input to the QGG algorithm. From top to bottom, the signals are from microphones near the mouth, near the forehead, near the nose, and near the throat, and (bottom) the electroglottogram. This section of speech corresponds to the “t” in the midst of “...what is...”. The throat microphone signal, **T**, has a noticeably smaller ratio of high-frequency power near 3.45 s to power in the voiced regions (edges) than the other microphones. That is because the high frequencies get to **T** only from the mouth, whereas the low frequencies also come through the throat wall.

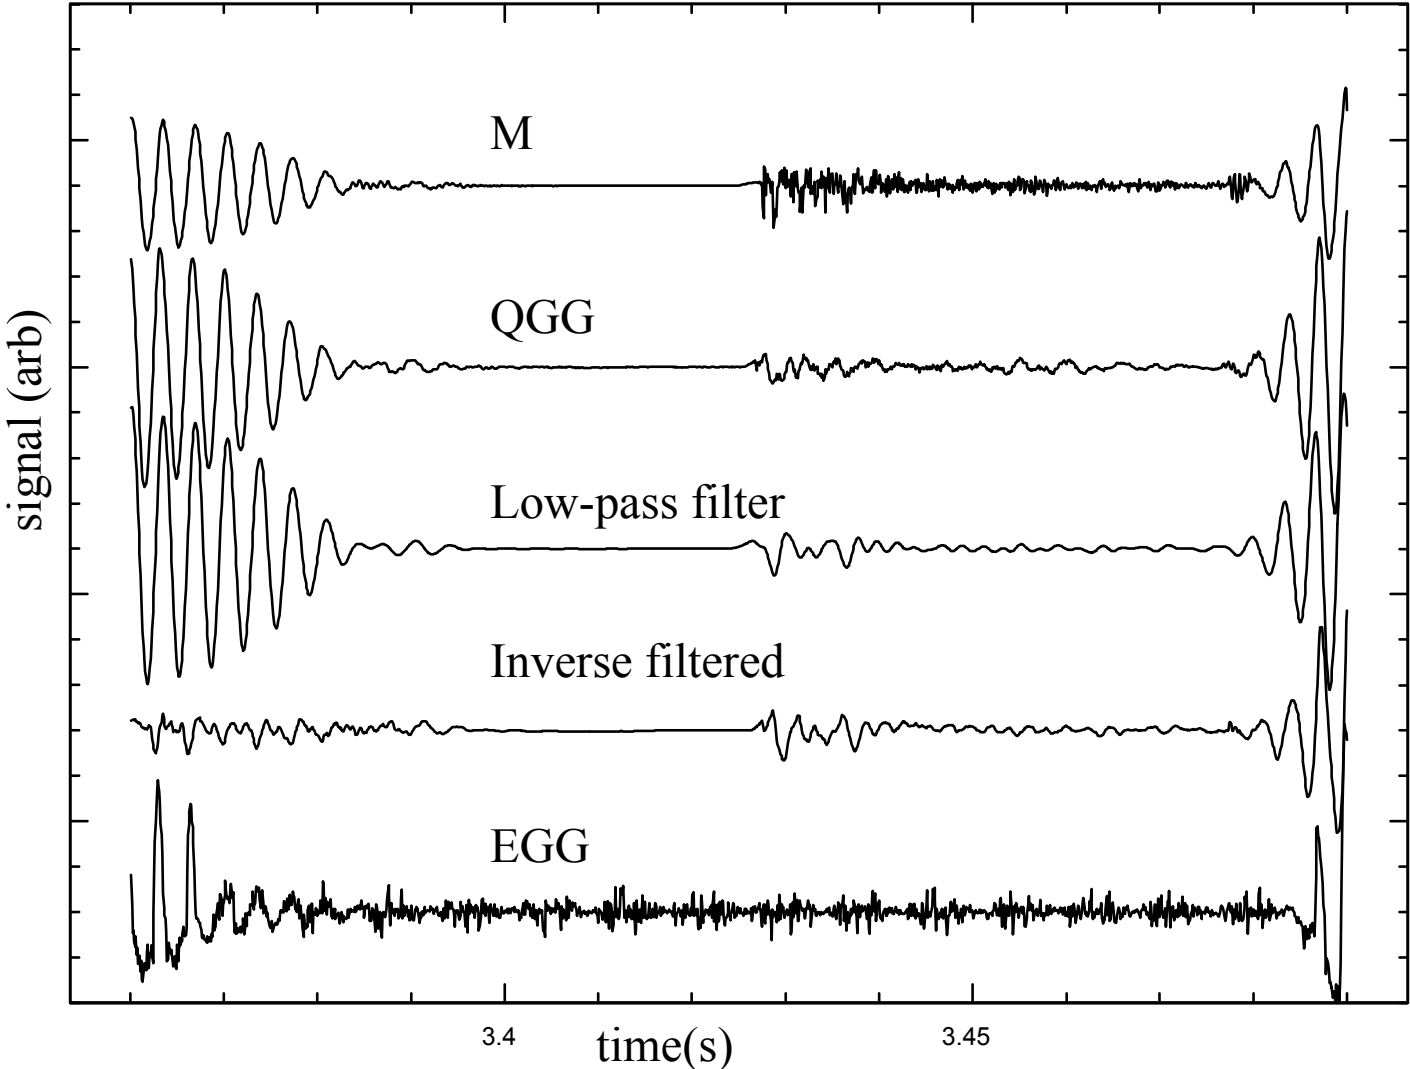


Figure 4. Output of the QGG algorithm, compared to other standard signals. From top to bottom, the signals are from the microphone near the mouth (M), the QGG, a low pass filter set just above peak f_0 , a standard inverse-filtering algorithm, and the EGG signal. This section of speech corresponds to the “t” in the midst of “...what is...”, and matches Figure 3.

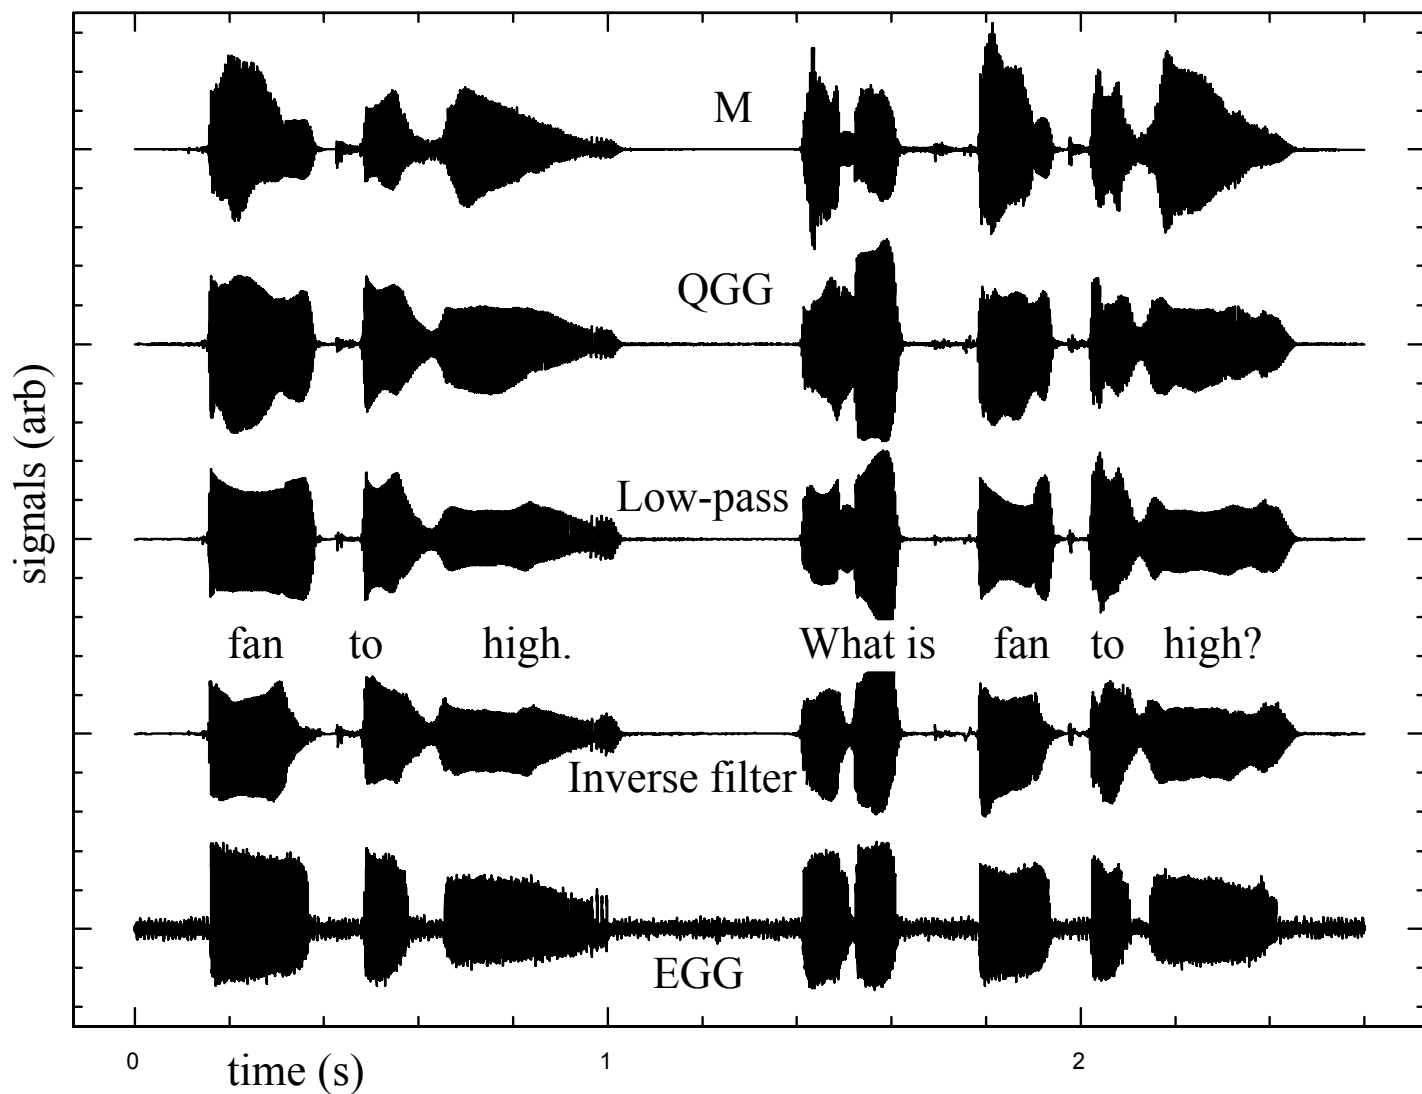


Figure 5. Output of the QGG algorithm, compared to other standard signals. From top to bottom, the signals are from the microphone near the mouth (M), the QGG, a low pass filter set just above peak f_0 , a standard inverse-filtering algorithm, and the EGG signal. The text is marked on the figure, approximately aligned with the audio. (The pause at 1.2 s has been shortened for a better visual display.)

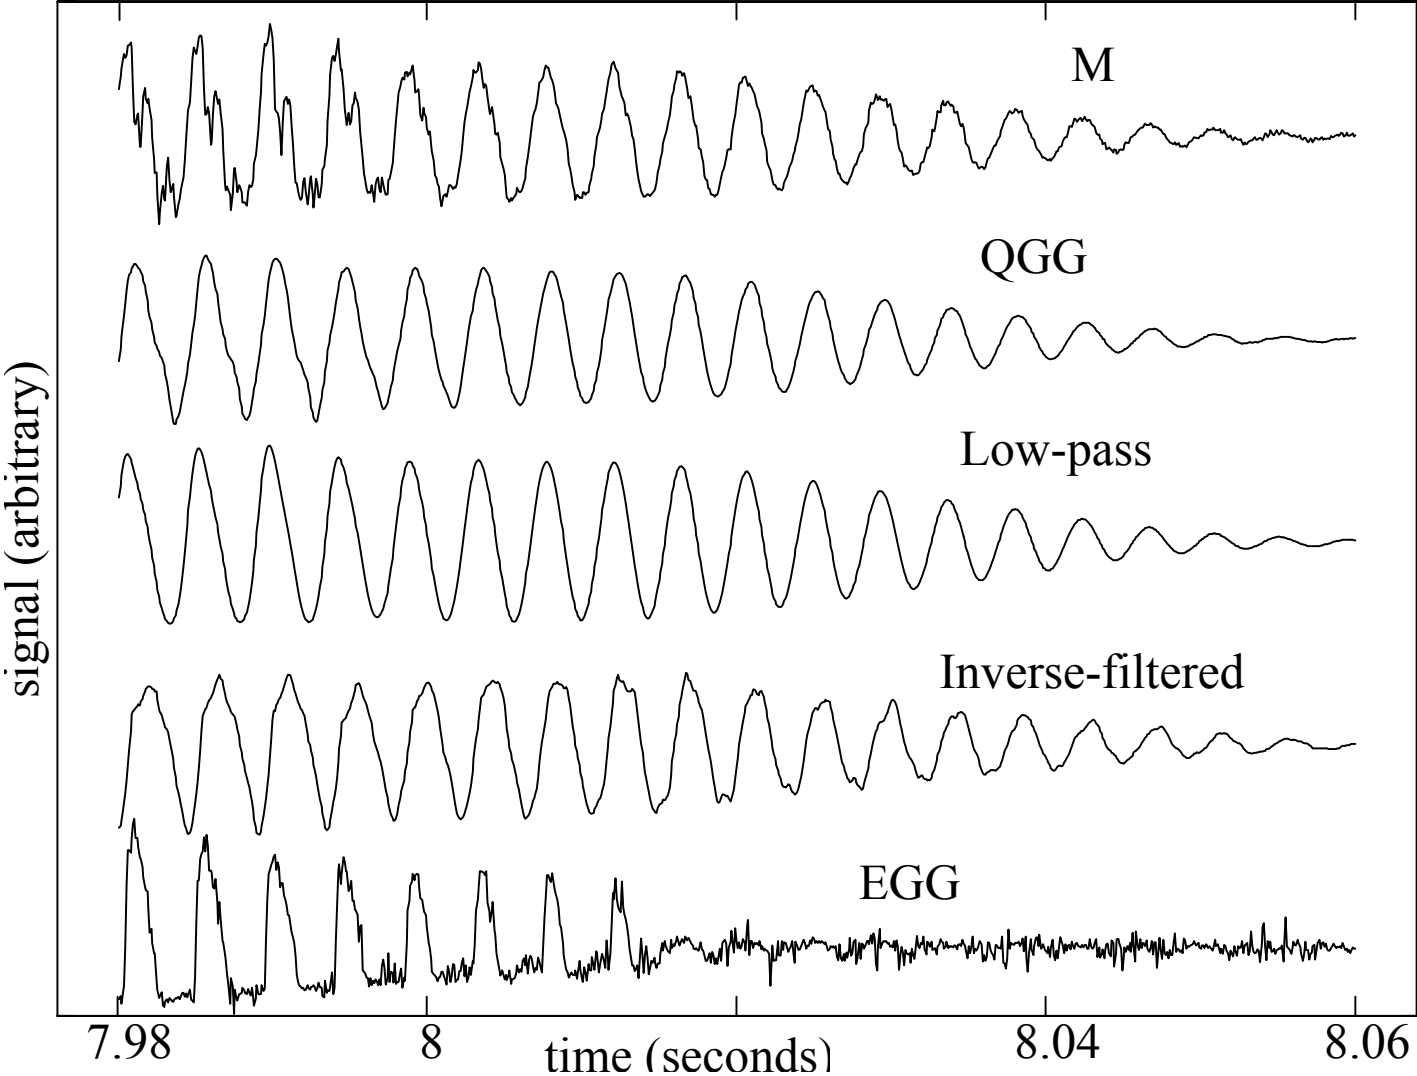


Figure 6. Comparisons of estimators of the glottal waveform. The signal is the off-glide [aɪ] in “high”. **M** is the raw signal from the microphone near the mouth, while “Low-pass” and “Inverse-filtered” are processed versions of **M**. QGG is derived from a microphone array, and EGG is the electroglottogram. The EGG stops more than 30 ms before the end of the glottal oscillation.




Figure 7. DFT spectra (16 ms windows) on either side of 8.015 s in Figure 6. The thick, black curve is before 8.015s (the window ends at 8.015s), when the glottal folds are colliding, and the wide, gray curve is after (window begins at 8.015s). The fundamental (225Hz) is essentially unchanged in amplitude, but the power in the harmonics drops by about 10 dB.

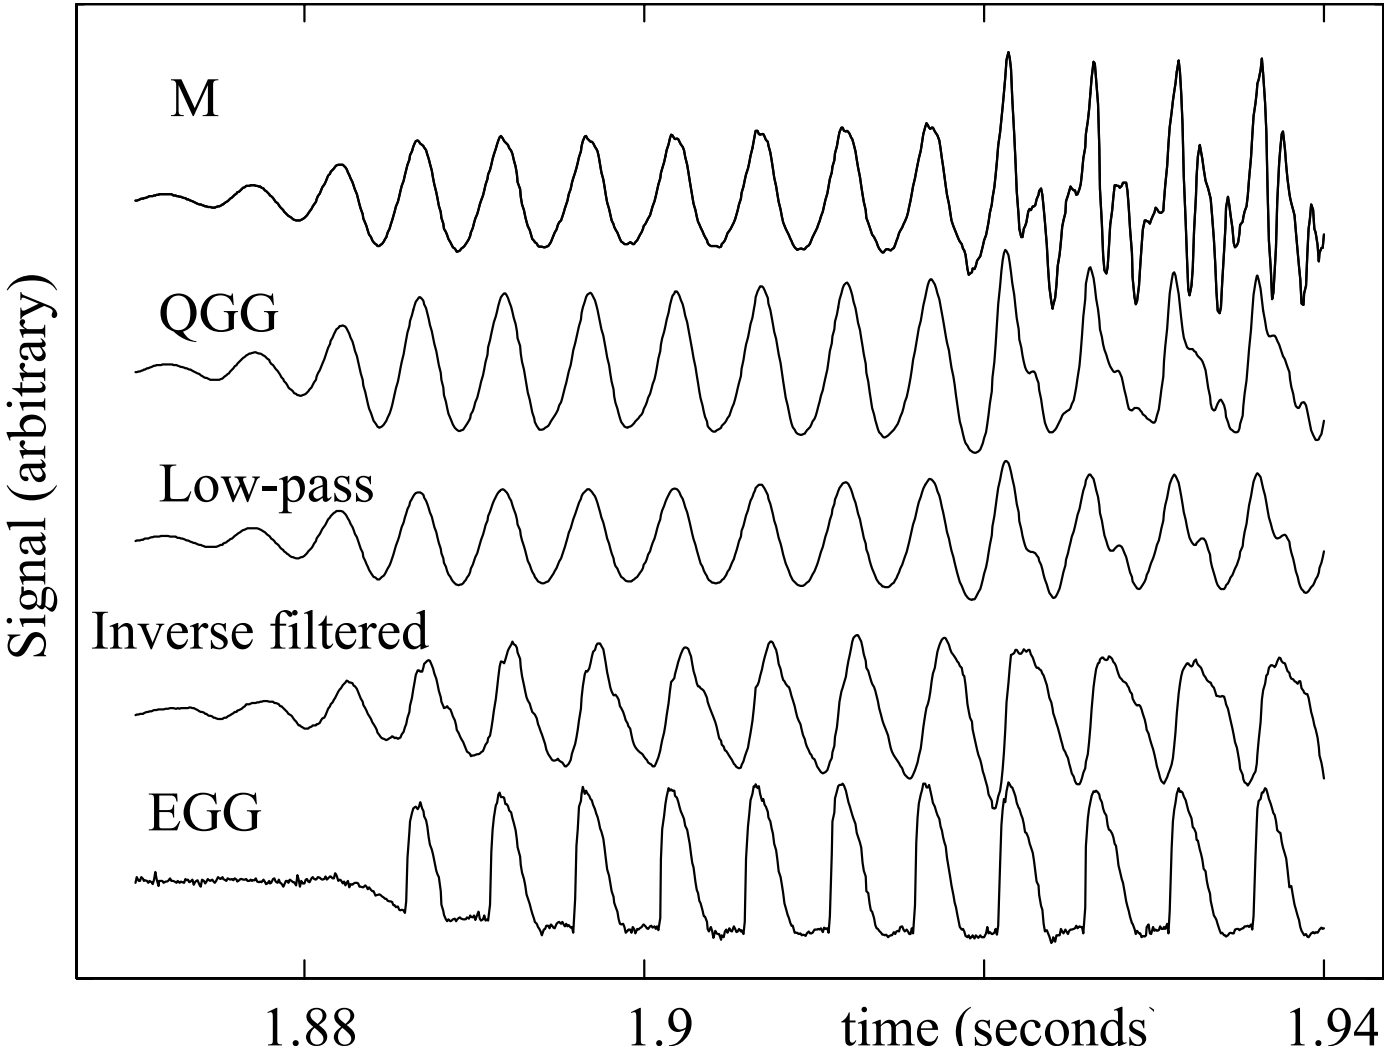


Figure 8. Beginning of the word ‘mosey’, at normal amplitude. Note that the EGG signal starts late. **M** is the raw signal from the microphone near the mouth, while “Low-pass” and “Inverse-filtered” are processed versions of **M**. QGG is derived from a microphone array, and EGG is the electroglottogram.

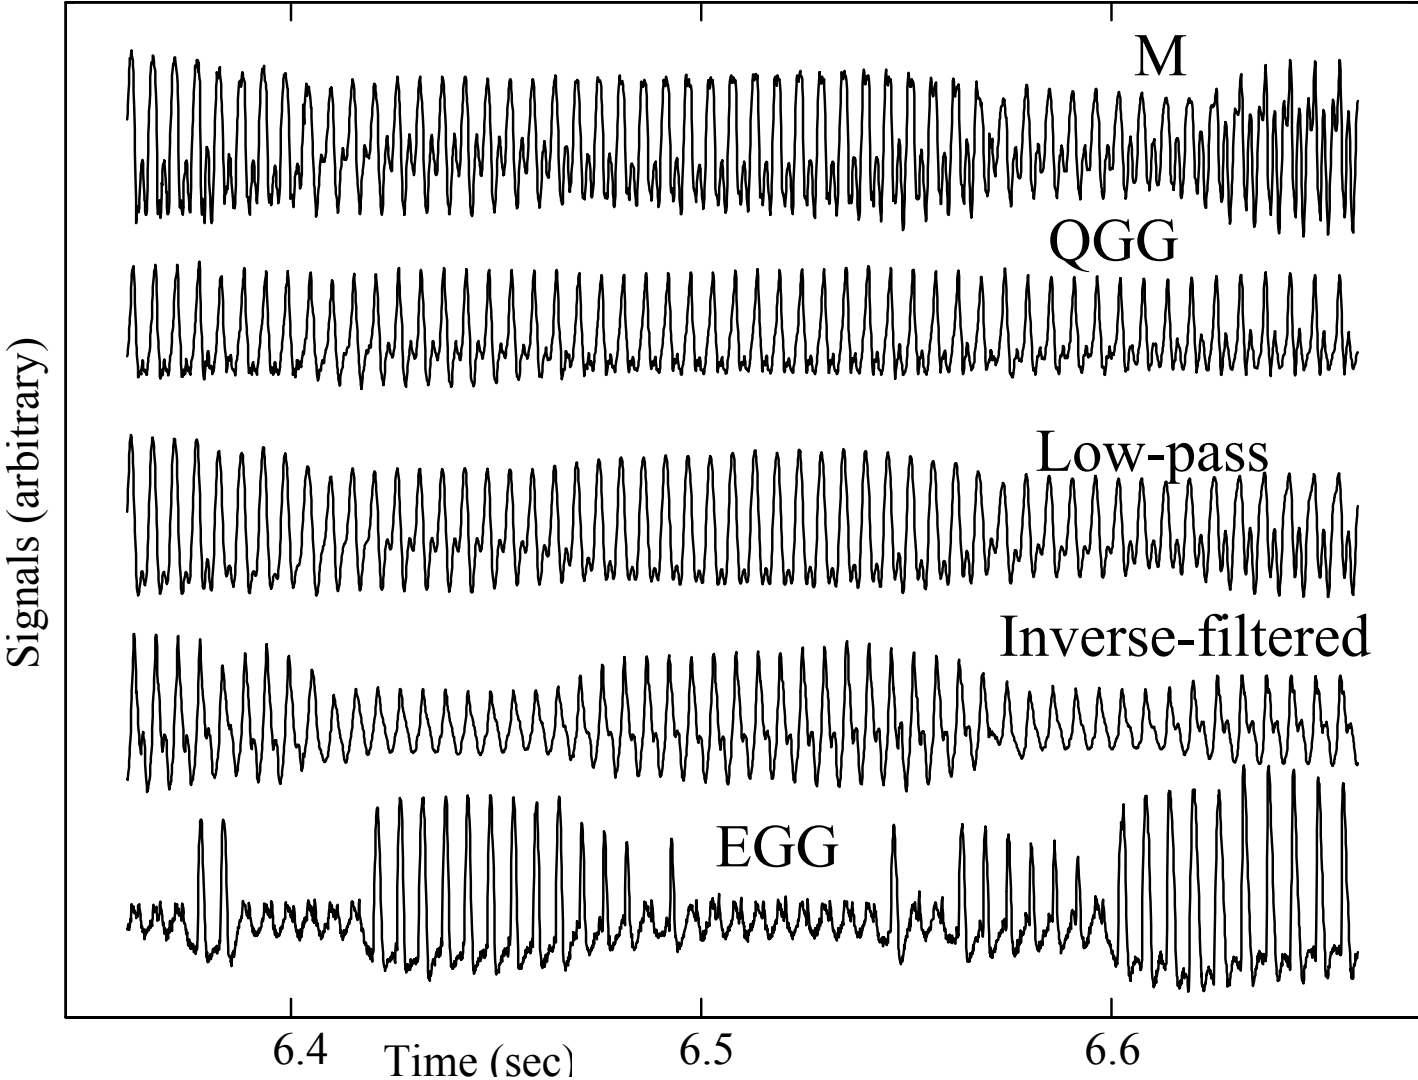


Figure 9. The middle of sustained low amplitude 'o' phonation showing major EGG changes (bottom) without large changes in the speech signal (top). The various signals are labeled.




Figure 10. Amplitude measurement (from QGG signal) vs. mean pitch of a syllable. Syllables with tone 1 (used in column X' of Table I and Table II) are shown as 'x'. Other tones are shown as dots. In this plot, the pitch-dependencies (from one tone to another) are larger than the dependence on the syllable, within the tone 1 syllables.

References

- ¹ Sondhi, M. M., 1975, "Measurement of the Glottal Waveform," *J. Acoustical Soc. Am.*, 57, 228-232.
- ² Baken, R. J. and Orlikoff, R. F., 2000, *Clinical Measurement of Speech and Voice*, Singular Publishing Group, ISBN 1-56593-869-0, p. 431-432 and references therein.
- ³ Baken, R. J. and Orlikoff, R. F., *op. cit.*, p.311-322 and references therein.
- ⁴ Farnsworth, D. W., 1940. "High-speed Motion Pictures of the Human Vocal Cords," *Bell Laboratories Record*, 18, pp. 203-208.
- ⁵ Baken, R. J. and Orlikoff, R. F., *op. cit.*, pp. 394-406 and references therein.
- ⁶ Sonneson, B. (1960). "On the anatomy and vibratory pattern of human vocal cords with special reference to a photo-electrical method for studying the vibratory movements," *Acta Oto-Laryngologica*, Supplementum **156**, 1-81.
- ⁷ Gerratt, B. R., Hanson, D. G., Berke, G. S., and Precoda, K. (1989). Photoglottography: A clinical synopsis, *J. Voice*, **5**(2), 98-105.
- ⁸ Baken and Orlikoff, *op. cit.*, pp. 322-324 and references therein.
- ⁹ Baken and Orlikoff, *op. cit.*, pp. 413-427 and references therein.
- ¹⁰ Fabre, P. 1957. "Un procédé électrique percutané d'inscription de l'accolement glottique au cours de la phonation: glottography de haute fréquence. Premiers résultats," *Bull. Acad. Natl. Med.* **141**, 66-69.
- ¹¹ Orlikoff, R. F., 1991, "Assessment of the Dynamics of Vocal Fold Contact from the Electroglottogram: Data from Normal Make Subjects", *J. Speech and Hearing Research*, **34**, 1066-1072.
- ¹² Sonesson, B. (1960) "On the anatomy and vibratory pattern of the human vocal folds," *Acta Otolaryngologica*, **Suppl. 156**, 1-80.
- ¹³ Timcke, R., von Leden, H., and Moore, P. (1958). "Laryngeal vibrations: Measurements of the glottic wave: Part 1. The normal vibratory cycle," *Archives of Otolaryngology*, **68**, 1-19.
- ¹⁴ Baer, T., Löfqvist, A., and McGarr, N. S. (1983). "Laryngeal vibrations: A comparison between high speed filming and glottographic techniques," *J. Acoustical Soc. America*, **73**, 1304-1308.
- ¹⁵ Baer, T., Titze, I. R., and Yoshioka, H. (1983). "Multiple simultaneous measures of vocal fold activity," in D. M. Bless and J. M. Abbs (Eds.), *Vocal Fold Physiology: Contemporary research and clinical issues* (pp. 229-237). San Diego, CA: College Hill Press.
- ¹⁶ Gilbert, H. R., Potter, C. R., and Hoodin, R. (1984). "Laryngograph and a measure of vocal fold contact area," *J. Speech and Hearing Research*, **27**, 556-565.
- ¹⁷ Baken and Orlikoff, *op. cit.*, pp. 433-436 and references therein.
- ¹⁸ Ishizaka, K., Matsudaira, M., and Kaneko, T, 1976. "Input acoustic-impedance measurement of the subglottal system," *J. Acoust. Soc. Am.* 60:1, pp. 190-196.
- ¹⁹ Milutinović, Z., Mijić, M., Djurica, S. (1997). "Activity of the subglottic voice ('chest resonator'): an echo-tomographic and acoustic study," *Eur. Arch. Otorhinolaryngol* 254: 292-297.
- ²⁰ K. Stevens (1998), *Acoustic Phonetics*, MIT Press, ISBN 0-262-19404-X. pp. 66-68, 76, 94.
- ²¹ Laryngograph, Ltd. 1 Foundry Mews, Tolmers Square, London NW1 2PE, U.K.
- ²² Entropic Research Laboratory. Code attributed to Brian Sublett, John Shore, David Talkin, and Derek Lin. It does a 14th order LPC analysis with a Hanning window, with a pre-emphasis coefficient of 0.97. Now available at <http://www.speech.kth.se/wavesurfer>.
- ²³ K. Stevens (1998), *op. cit.*, p.200-202.
- ²⁴ Weider, Sol, 1973. *The Foundations of Quantum Theory*, Academic Press, New York, pp. 44-59.
- ²⁵ Fant, G. (1962). "Formant Bandwidth Data". *Speech Transmission Laboratory Quarterly Progress and Status Report 2-3*, Royal Institute of Technology, Stockholm, 1-3.
- ²⁶ Fujimura, O. and Lindqvist, J. (1971). *Sweep-tone measurements of vocal tract characteristics*. J. Acoustical Society of America, **49**, 541-558.
- ²⁷ K. Stevens *op. cit.* p. 152-167, 264
- ²⁸ K. Stevens, *op. cit.*, p. 196-198 and references therein.
- ²⁹ Press, W. H., Teukolsky, S., Vetterling, W. T., Flannery, B. P., *Numerical Recipes in C: The Art of Scientific Computing*, 2nd edition, Cambridge University Press, 1992, ISBN 0-521-43108-5. p. 564-565.
- ³⁰ Colton, R. H., Conture, E. G., (1990). "Problems and pitfalls of electroglottography," *J. Voice* **4**, pp. 10-24.
- ³¹ Hanson, D. G., Gerratt, B. R., Berke, G. S., (1990). "Frequency, intensity, and target matching effects on photoglottic measures of open quotient and speed quotient," *J. Speech Hearing Res.* **33**, pp. 45-50.
- ³² Titze, I. R. (1988). "The physics of small amplitude oscillation of the vocal folds," *J. Acoust. Soc. Am.* **83**(4), 1536-1552.
- ³³ K. Stevens, *op. cit.* p. 91.
- ³⁴ Ishizaka, K. and Matsudaira, M. (1968). "What makes the vocal cords vibrate?," In *Proceedings of Sixth International Congress of Acoustics*, Tokyo, B1-3.

- ³⁵ Ishizaka, K. and Flanagan, J. L. (1972). "Synthesis of voiced sounds from a two-mass model." *Bell System Technical Journal*, **51**, 1233-1268.
- ³⁶ Fry, D. B., (1955), "Duration and intensity as physical correlates of linguistic stress," *J. Acoustical Soc. Am.* **30**, 765-769
- ³⁷ Fry, D. B. (1958), "Experiments in the perception of stress," *Language and Speech*, **1**, 126-152.
- ³⁸ Bolinger, D. L. (1958), "A theory of pitch accent in English," *Word*, **14**, 109-149.
- ³⁹ Lieberman, P. (1960), "Some acoustic correlates of word stress in American-English," *J. Acoustic Soc. Am.*, **32**, 451-454.
- ⁴⁰ Hadding-Koch, K. (1961), *Acoustico-Phonetic Studies in the Intonation of Southern Swedish*, C. W. K. Gleerup, Lund, Sweden.
- ⁴¹ Pollock, K. E., Brammer, D. M. and Hageman, C. F. (1990), "An acoustic analysis of young children's productions of word stress", *J. Phonetics*, **21**, 183-203.
- ⁴² Kehoe, M., Stoel-Gammon, C., and Buder, E. H. (1995), "Acoustic Correlates of Stress in Young Children's Speech", *J. Speech and Hearing Research*, **38**, 338-350.
- ⁴³ Sereno, J. A., Jongman, A. (1995), "Acoustic Correlates of Grammatical Class," *Language and Speech* **38(1)**, 57-76.
- ⁴⁴ Beckman, M. (1986), *Stress and non-stress accent*, Dordrecht: Foris Publications.
- ⁴⁵ Sundberg, J., Elliot, N., Gramming, P., and Nord, L. (1993) "Short-term variation of subglottal pressure for expressive purposes in singing and stage speech – a preliminary investigation," *J. Voice* **7(3)** 227-234.