

(51) International Patent Classification:
G10L 15/02 (2006.01)(21) International Application Number:
PCT/US2012/062549(22) International Filing Date:
30 October 2012 (30.10.2012)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant: **NUANCE COMMUNICATIONS, INC.**
[US/US]; One Wayside Road, Burlington, MA 01803-4613 (US).(72) Inventors: **KRINI, Mohamed**; Elchinger Weg 10, 89075 Ulm (DE). **SCHALK-SCHUPP, Ingo**; Marktplatz 30, 89312 Günzburg (DE). **HANNON, Patrick**; Wagnerstr. 60, 89077 Ulm (DE).(74) Agents: **MURPHY, Timothy, M.** et al.; Sunstein Kann Murphy & Timbers LLP, 125 Summer Street, Boston, MA 02110 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,


AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: SPEECH ENHANCEMENT



(57) **Abstract:** A speech signal processing system is described for use with automatic speech recognition and hands free speech communication. A signal pre-processor module transforms an input microphone signal into corresponding speech component signals. A noise suppression module applies noise reduction to the speech component signals to generate noise reduced speech component signals. A speech reconstruction module produces corresponding synthesized speech component signals for distorted speech component signals. A signal combination block adaptively combines the noise reduced speech component signals and the synthesized speech component signals based on signal to noise conditions to generate enhanced speech component signals for automatic speech recognition and hands free speech communication.



Speech Enhancement

TECHNICAL FIELD

[0001] The present invention relates to noise reduction in speech signal processing for automatic speech recognition and hands free speech communication.

BACKGROUND ART

[0002] In speech signal processing for automatic speech recognition (ASR) and hands free speech communication, a microphone signal is usually first segmented into overlapping blocks of appropriate size and a window function is applied. The speech signal processing can be performed in the time domain and/or in the frequency domain. Time domain processing operates directly on speech waveform signals while frequency domain processing operates on spectral representations of the speech signal.

[0003] Operations carried out in the frequency domain are achieved by a short-term Fourier transform (STFT). In this process, the sequence of sampled amplitude values x in dependence of the sample index i is multiplied with a window sequence w for every R th sample and then discretely transformed to the frequency domain. This step is called analysis, and its realization is often referred to as an analysis filter bank:

$$X(k, \mu) = \sum_{i=0}^{L-1} x(i + Rk) w(i) e^{-j \frac{2\pi \mu i}{N}}$$

for each sample i , frame k , frequency bin μ , frame shift R , window length L , and DFT size N . After processing in the frequency domain, the resulting spectrum is transformed to the time domain again by an inverse STFT. In analogy to the previous step, this one is called synthesis, and its implementation is referred to as the synthesis filter bank.

[0004] Frequency domain processing produces noisy short-term spectra signals. In order to reduce the undesirable noise components while keeping the speech signal as natural as possible, SNR-dependent (SNR: signal-to-noise ratio) weighting coefficients are computed and applied to the spectra signals. Common noise reduction algorithms make assumptions

to the type of noise present in a noisy signal. The Wiener filter for example introduces the mean of squared errors (MSE) cost function as an objective distance measure to optimally minimize the distance between the desired and the filtered signal. The MSE however does not account for human perception of signal quality. Also, filtering algorithms are usually applied to each of the frequency bins independently. Thus, all types of signals are treated equally. This allows for good noise reduction performance under many different circumstances.

[0005] However, mobile communication situations in an automobile environment are special in that they contain speech as their desired signal. The noise present while driving is mainly characterized by increasing noise levels with lower frequency. Speech signal processing starts with an input audio signal from a speech-sensing microphone. The microphone signal represents a composite of multiple different sound sources. Except for the speech component, all of the other sound source components in the microphone signal act as undesirable noise that complicates the processing of the speech component.

[0006] Separating the desired speech component from the noise components has been especially difficult in moderate to high noise settings, especially within the cabin of an automobile traveling at highway speeds, when multiple persons are simultaneously speaking, or in the presence of audio content. Often in high noise conditions, only a low output quality is achieved. Usually, speech signal components are heavily distorted to such an extent that desired speech components are masked by the background noise. Standard noise suppression rules classify these parts as noise; as a consequence, a maximum attenuation is applied.

SUMMARY

[0007] Embodiments of the present invention are directed to an arrangement for speech signal processing for automatic speech recognition and hands free speech communication. The processing may be accomplished on a speech signal prior to speech recognition, for example, with mobile telephony signals and more specifically in automotive environments that are noisy.

[0008] A signal pre-processor module transforms an input microphone signal into corresponding speech component signals. A noise suppression module applies noise reduction to the speech component signals to generate noise reduced speech component signals. A speech reconstruction module produces corresponding synthesized speech component signals for distorted speech component signals. A signal combination block adaptively combines the noise reduced speech component signals and the synthesized speech component signals based on signal to noise conditions to generate enhanced speech component signals for automatic speech recognition and hands free speech communication.

[0009] In further specific embodiments, the speech component signals may be time-domain speech component signals or frequency-domain speech component signals. The speech reconstruction module may use a non-linear function of speech components for producing the synthesized speech component signals. For example, distorted or missing harmonics of the speech component signals can be restored by the non-linear function.

[0010] The speech reconstruction module may use relevant features extracted from the distorted speech component signals for producing the synthesized speech component signals. The speech reconstruction module may use a source-filter model for producing the synthesized speech component signals. The speech reconstruction module may use a parametric model of envelope shape for producing the synthesized speech component signals. The speech reconstruction module may include a noise pre-processing module for noise suppression of the distorted speech component signals to provide speech synthesis reference signals. The speech reconstruction module may divide the distorted speech component signals into different frequency bands for producing the synthesized speech component signals.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] Figure 1 shows various functional blocks in a system for speech signal processing for automatic speech recognition and hands free speech communication that uses synthesized speech reconstruction according to an embodiment of the present

invention.

[0012] Figure 2 shows various steps in a method of processing speech according to an embodiment of the present invention.

[0013] Figure 3 A-B shows time frequency graphs of a noisy speech signal of a conventionally noise reduced speech signal without and with speech reconstruction..

DETAILED DESCRIPTION

[0014] Embodiments of the present invention are directed to enhancing highly distorted speech using a computationally efficient partial speech signal reconstruction algorithm. The partial signal reconstruction exploits the correlation of the speech components and replaces highly disturbed speech component signals with a corresponding synthesized speech component signal. The synthesized speech component signal is based on speech components with sufficient signal-to-noise ratio (SNR) using a nonlinear operator (nonlinear characteristic) as well as relevant features extracted from the noisy speech signal. The synthesized speech component signals are combined with the noise reduced speech component signals to generate an enhanced speech component signal output. Specific embodiments may operate on speech component signals in either the time domain or in the frequency domain. The following discussion is specific to operation in the frequency domain, but those in the field will appreciate that the principle of the invention is equally applicable to signal processing of time domain speech component signals.

[0015] Figure 1 shows various functional blocks in a system and Figure 2 shows various steps in a method for speech signal processing for automatic speech recognition and hands free speech communication that uses synthesized speech reconstruction according to an embodiment of the present invention. Initially, an analysis filter bank **104** in noise control module **101** transforms an input microphone signal x into corresponding short term spectra signals X , step **201**. A noise suppression module **105** then applies noise reduction to the short term spectra signals X to generate noise reduced spectra signals $\hat{S}_{NR}$, step **202**. A speech reconstruction module **102** produces corresponding synthesized spectra signals

$\hat{S}_{est}$ for distorted short term spectra signals X , step **203**. And a signal combination block **103** includes an adaptive mixer **106** that adaptively combines the noise reduced spectra signals $\hat{S}_{NR}$ and the synthesized spectra signals $\hat{S}_{est}$ based on signal to noise conditions and extracted features from the input (e.g. voiced/unvoiced classification, variability of the short-term spectrum), step **204**, and transforms the resulting enhanced speech signal $\hat{S}_{enhanced}$ in a synthesis filter bank **112** performing an inverse STFT and summing the frames back into a time-domain enhanced speech signal $\hat{S}_{enhanced}$ for automatic speech recognition and hands free speech communication, step **205**.

[0016] Looking in greater detail at the speech reconstruction module **102**, a noise pre-processing module **107** utilizes a noise suppression algorithm to attenuate undesired signal components. Typically, high noise attenuation (e.g., about 30 dB) is applied by the pre-processing module **107** to generate a greatly noise-reduced reference signal $\hat{S}_{pp}$ for properly regenerating distorted speech components.

[0017] Non-linear operator module **108** regenerates distorted or missing harmonics of the noise-reduced reference signal $\hat{S}_{NR}$ utilizing a nonlinear function that applies a nonlinear characteristic to a harmonic signal to produce a non-linear speech signal S_{nl} that has sub- and super-harmonics. Normally a nonlinear characteristic can only be applied efficiently (in terms of computational complexity) to speech signals in the time-domain, and using a nonlinear function for speech signal processing in the frequency domain comes at a much higher computational cost. So for efficiency reasons, in the frequency domain the non-linear operator module **108** may apply the non-linear operation using auto-convolution of a subset of the current sub-band signals. For example, the non-linear operator module **108** may divide the short-term spectrum into different frequency ranges (e.g., into 1 kHz bands from 100 Hz up to 4 kHz) and depending on the measured variability within each band, an upper frequency limit for the convolution can be adaptively determined. In embodiments operating on time domain speech component signals, the non-linear operator module **108** may apply a non-linear operation based on, for example, a quadratic characteristic or a half-way rectification of the speech component

signals.

[0018] An excitation signal generator 109 receives the non-linear speech signal S_{nl} and for voiced signals generates a synthetic excitation signal $\tilde{S}_{exc}$ having harmonics at the desired frequencies but with biased amplitudes. In order to reduce the bias, the excitation signal generator 109 may divide the incoming non-linear speech signal S_{nl} by its envelope.

[0019] A separate envelope estimation module 110 estimates the frequency-domain envelope $\bar{S}_{env}$ of the original short term spectra signals X . The envelope estimation module 110 can operate on-line in real time using a novel approach having a very low computational power cost. The envelope estimation module 110 analyzes the filter coefficients of the noise reduced spectra signals $\hat{S}_{NR}$ from the noise suppression module 105 to distinguish frequency bins of the preliminary envelope with high and low speech signal-to-noise power estimation. The high SNR bins are used as supporting points for the re-estimation of the bins with low SNR.

[0020] Synthesis combination module 111 uses a source-filter model to combine the estimated envelope $\bar{S}_{env}$ and the synthesis excitation signal $\tilde{S}_{exc}$ to generate synthesized spectra signals $\hat{S}_{est}$ that reconstruct the original undisturbed speech spectrum as accurately as possible.

[0021] The adaptive mixer 106 adaptively combines the synthesized spectra signals $\hat{S}_{est}$ with the noise reduced spectra signals $\hat{S}_{NR}$ to produce enhanced frequency domain speech signals $\hat{S}_{enhanced}$. At low SNRs, highly distorted noise reduced spectra signals $\hat{S}_{NR}$ are replaced by the synthesized spectra signals $\hat{S}_{est}$. At more moderate SNR conditions (i.e. medium and high SNRs), only the magnitudes of the spectra signals $\hat{S}_{est}$ and of the noise reduced spectra signals $\hat{S}_{NR}$ are adaptively combined. The adaptive mixer 106 afterwards combines the resulting magnitude estimate with the phase of the noise reduced spectra

signals $\hat{S}_{NR}$. Thus speech reconstruction is performed just for degraded voiced speech components. The synthesized phase is only applied for highly degraded voiced speech components.

[0022] A voice activity detector can be used to detect the presence of speech segments, and a voiced speech detector can be used to detect voice speech components. To detect corrupted voiced speech components, the adaptive mixer 106 can perform a comparison at each frequency bin between the synthesized spectra signals $\hat{S}_{est}$ and the noise reduced spectra signals $\hat{S}_{NR}$. The amplitudes of the synthesized spectra signals $\hat{S}_{est}$ at the pitch frequency and at harmonics are always greater than the noise reduced spectra signals $\hat{S}_{NR}$ (for degraded speech components [highly degraded harmonics]. If the harmonics are degraded the noise reduction filter will apply a maximum attenuation.) Thus the speech reconstruction is only performed at those frequency bins where the synthesized spectra signals $\hat{S}_{est}$ are greater than the noise reduced spectra signals $\hat{S}_{NR}$ by a specified margin.

[0023] Figure 3 A-B shows time-frequency analyses of a noisy speech signal measured in a car at high speed (at about 160 km/h). Fig. 3A shows a time-frequency analysis of a conventional noise reduced signal. Fig. 3B shows an analysis of the outcome after partial speech reconstruction. The black ellipses in Fig. 3B denote the reconstructed speech components. Comparing these results one can see that missing speech components are successfully regenerated after using partial reconstruction.

[0024] Looking back at the envelope estimation more formally, the task of the envelope estimation module 110 is to estimate the smoothed spectrum of the clean speech signal as closely as possible. A decision is made as to how to separate the contributions of the voiced excitation and the vocal tract to the clean speech spectrum so as to supply comparable results. A time-smoothed version $\hat{\Phi}_{\hat{D}_T}$ of the noisy speech amplitude spectrum $\hat{\Phi}_X$ can be used as an initial estimate for the noise spectrum, and the smoothing can be implemented using a first-order infinite impulse response (IIR) filter in forward and

backward direction with a discrete-time smoothing constant $\gamma_{\bar{i}}$:

$$\begin{aligned}\hat{\Phi}_{\hat{D}_{\bar{i}}}(k=0, \mu) &:= \hat{\Phi}_X(k=0, \mu) \\ \hat{\Phi}_{\hat{D}_{\bar{i}}}(k, \mu) &:= \gamma_{\bar{i}} \hat{\Phi}_{\hat{D}_{\bar{i}}}(k-1, \mu) + (1-\gamma_{\bar{i}}) \hat{\Phi}_X(k, \mu)\end{aligned}$$

where the discrete-time smoothing constant $\gamma_{\bar{i}}$ is given by:

$$\gamma_{\bar{i}} = \exp\left(\frac{-R}{\tau_{\bar{i}} f_s}\right)$$

where R is the frame shift, $\tau_{\bar{i}}$ is a selected (continuous) decay time constant, and f_s is the sampling frequency. Note that the smoothing can also be applied in the linear domain to reduce computational cost.

[0025] The smoothed amplitude spectrum of the microphone or of the noise reduced signal is a good estimator of the envelope at high SNR conditions. But at low SNR levels, reliable estimation no longer can be achieved, especially at low frequencies if speech formants are masked by the background noise. Hence the envelope estimation module 110 can develop the envelope estimation using a parametric model for the shape of the envelope $\bar{S}_{env}$. To extrapolate towards lower frequencies, a logarithmic parabola can be used, while for interpolating between higher frequencies, a logarithmic hyperbola form can be applied. These two specific shapes have been heuristically derived from observed envelope shapes and then modified by several parameters. The hyperbola shape is fully determined by its maximum steepness, while the parabola shape can use three parameters: (1) the starting slope at the good bin with the lowest frequency, (2) the amount of curvature, and (3) the maximum slope allowed. For that purpose, the envelope estimation module 110 can easily extract several features from the noise reduced spectra signals $\hat{S}_{NR}$ to determine a useful estimate for the curves' parameters. A functional dependency between the features and the optimal parameters can be established off-line from a training data set. This way, the cost of the computational power under operating conditions stays minimal. The features used may be, for example, the slope of the reference spectrum for the parabola shape, and the distance between the neighboring good bins for the hyperbola shape. Applying the resulting parametric estimation process on a test data set yields better estimation results than a state-of-the-art codebook approach which is also much more

expensive in terms of memory and computational power requirements.

[0026] For off-line parameter optimization based on a training data set of clean speech, the same definition can be used for the known signal's amplitude spectrum $\hat{\Phi}_S$:

$$\bar{S}(k, \mu) := \text{IIR} \left\{ 10 \log_{10} \hat{\Phi}_S(k, \mu) \right\}$$

The estimation task is the minimization of the sum of absolute differences in each relevant frame.

[0027] On modern computers, unsupervised vector quantization methods such as a codebook approach are mainly used to reduce the dimension of a feature space in order to save memory space and computation power. But codebook based estimation is still in common use in embedded systems, where these resources are limited. In the training step, codebooks store information on prototype feature vectors extracted from a database. In the testing step, incomplete feature vectors can be matched against all codebook entries by use of a cost or distance function. The codebook entry that best matches the incomplete vector is used to fill in missing features. In the task at hand, missing features are frequency bins where the corresponding filter coefficients lie below a certain threshold, thus indicating a low SNR. In these bins, the estimate amplitude value is dominated by the noise power, and the speech power is bound to be much lower. A codebook can be implemented using the well-known Linde-Buzo-Gray (LBG) algorithm for training on clean speech signal envelopes using 256 feature vectors each with M amplitude values.

[0028] Besides the codebook approach, other estimation methods can be considered, for example, based on bin-wise computation of available spectral information. In frequency bins where the noise estimation is accurate and close to 0, the noise reduction filter will output coefficients near 1 or 0 dB on a logarithmic scale. Here, the microphone spectrum provides a fairly good estimate for the speech spectrum at high SNR levels.

[0029] *Microphone.* The noisy microphone signal's amplitude spectrum $\hat{\Phi}_X$ can be smoothed and used as a first approximation:

$$\bar{S}_{\text{Mic}}(k, \mu) := \bar{X}(k, \mu)$$

As is to be expected, this approximation only works well where no or very little noise is present. It can however be useful as a source of information applied by other estimators.

[0030] *Average.* While similar to the microphone estimator, an average estimation method indirectly takes into account the filter coefficients $H(k, \mu)$ available from a preceding noise reduction. These can range from the spectral floor up to a value of 1 and contain information on how much noise is present in each bin μ and frame k . Higher values indicate a better SNR. The actual implementation can rely on the smoothed amplitudes of the noise reduced spectra signals $\hat{S}_{NR}$:

$$\bar{S}_{\text{Avg}}(k, \mu) := \frac{1}{2} \left(\bar{X}(k, \mu) + \hat{S}_{NR}(k, \mu) \right)$$

[0031] As explained above, in a separate processing stage, the frequency-domain envelope of the original speech signal is estimated. The estimated envelope is then imprinted on the generated excitation signal in order to reconstruct the original speech spectrum as accurately as possible. The envelope estimation is implemented in a novel way that can be applied online at a very low computational power cost. First, the noise suppression filter's coefficients are analyzed to distinguish bins with good and bad speech signal power estimation. The well-estimated bins are then used as supporting points for the estimation of the badly estimated ones. The estimation itself consists in the use of a parametric model for the envelope shape. For extrapolating towards low frequencies, a logarithmic parabola is used. The shape is derived from observed envelope shapes and is modified by several parameters, such as the parabola's curvature. Several features can easily be extracted from the signal, such as the position of the lowest well estimated frequency bin. These features are used to determine a good estimate for the curves' parameters.

[0032] Disturbances in the low frequency spectrum are prevalent in hands-free automobile telecommunication systems. Typically a number of consecutive low-frequency bins do not offer enough speech signal information for estimating the envelope at that

frequency bin. In a linear extrapolation approach, missing envelope information can be approximated by constructing a straight line in the logarithmic amplitude spectrum.

[0033] The lowest well-estimated bin's index in a frame k is referred to as $\mu^{\circ}(k)$, and the logarithmic amplitude value in that place is $\hat{S}_{\text{Avg}}(k, \mu^{\circ}(k))$ (NB: other preliminary estimators could be used, for example, $\hat{S}_{\text{NR}}(k, \mu^{\circ}(k))$). The line can be fixed to this point retaining one degree of freedom, namely its slope $m_{\text{line}}(k)$. This parameter can be estimated from the signal's features with the estimated amplitude values calculated iteratively. This gives a straight line in logarithmic space depending on the slope $m_{\text{line}}(k)$, which is set to a value leading to a good approximation of the actual spectral amplitudes.

[0034] *Parabolic Extrapolation.* It is not known directly whether the low-frequency band contains an obscured formant or not. As an extension to the aforementioned linear extrapolation, a parabolic estimation model can be parameterized by the slope of the reference estimation to guess the optimal course of the envelope. This can be more flexible than the linear method, including it as a special case. Hence, it bears an equal-or-better potential performance but also introduces higher worst-case errors. From the reference estimation, a reference slope m_{R} can be computed:

$$m_{\text{R}} := \sqrt{\frac{\hat{S}_{\text{Avg}}(k, \mu^{\circ}(k) - 1)}{\hat{S}_{\text{Avg}}(k, \mu^{\circ}(k) + 1)}}$$

and used to determine the parabola's initial slope $m^{\circ}(k) = m(k, \mu^{\circ}) = m_{\text{R}}$ at μ° . The curve can be fixed by setting the remaining degree of freedom, called here the curvature $J(k)$. Additionally, a restriction on the curve's maximum slope m_{max} can be introduced and the minimum slope can be restricted to that of the estimated noise amplitude spectrum:

$$m_{\hat{D}}(k, \mu) := \frac{\hat{\Phi}_{\hat{D}}(k, \mu)}{\hat{\Phi}_{\hat{D}}(k, \mu + 1)}$$

[0035] For lower frequencies, $\hat{S}(k, \mu) = p(k, \mu)$, to be precise, from $\mu=0$ up to $\mu=\mu^{\circ}(k)$. For these frequencies we use a parametric model to achieve a reliable estimator for the envelope. For the remaining frequencies (i.e. from $\mu=\mu^{\circ}(k+1)$ up to $\mu=M-1$ one can use the

average envelope (i.e. $\hat{S}(k, \mu) = \hat{S}_{avg}(k, \mu)$). Given the previously mentioned parameters, the parabola $p(k, \mu)$ can be calculated iteratively by counting the bin index μ down from μ° with the following algorithm:

$$p(k, \mu) := J(k) \cdot p(k, \mu + 1) \cdot \max\{m_D, \min\{m_{\max}, m(k, \mu)\}\}$$

where the slope $m(k, \mu)$ is:

$$m(k, \mu) := \begin{cases} m^\circ(k) & \text{for } \mu = \mu^\circ, \\ \frac{p(k, \mu + 1)}{p(k, \mu + 2)} & \text{otherwise} \end{cases}$$

This results in a parabolic shape that is determined by the parameters initial slope $m^\circ(k)$, curvature $J(k)$, and minimum slope $m_{\min}(k)$, as well as on the course of the noise spectrum via m_D .

[0036] In the foregoing the construction of a parabolic curve depending on parameters was described. But it has not yet been formally defined how these parameters are derived from the features extracted from the available information. Functional dependencies between the features and the optimal parameters can be established off-line from a training data set to obtain a simple constant or linear dependency from a feature to a parameter. In this way, the computational power cost under operating conditions stays minimal.

[0037] Detecting functional dependencies can be performed in two steps. In the first one, a minimum search for the cost function can be performed on the available parameters. This establishes the optimum approximation that a model is capable of. The parameters defining this optimal approximation can be saved together with the feature extracted from information that would be available also in testing conditions. In the second step, the tuples of features and corresponding optimal parameters can be analyzed to detect approximate functional dependencies between them. If there appears to be a dependency, a zero- to second-order polynomial can be fit on the data to represent it. This can be evaluated in a testing step.

[0038] *Optimum Parameter Detection and Feature Extraction.* First, the testing data set can be searched for frames in which reconstruction can likely be successfully

performed. Criterion for exclusion may be insufficient voiced excitation power (with modified parameters). The band used goes from $f_{\text{low}} = 800$ Hz to $f_{\text{high}} = 2400$ Hz, and the threshold value is $P_{\text{VUD}} = 2$ dB. In the remaining frames, the noisy signals can be processed and the envelope estimation parameters for each gap in the smoothed spectrum can be optimized to approximate the known clean speech smoothed spectrum. The cost function C used in each frame k for finding the best fit of the parametric model $\hat{\bar{S}}(k, \mu)$ on the clean speech smoothed amplitude spectrum $\bar{S}(k, \mu)$ is the mean squared logarithmic spectral distortion:

$$C(k) := 20 \sum_{\mu=0}^{M-1} \left(\log_{10} \frac{\bar{S}(k, \mu)}{\hat{\bar{S}}(k, \mu)} \right)^2$$

[0039] Apart from the optimized parameters for each frame, several features can be extracted from the noisy speech signal. These features may be likely to carry valuable information on the corresponding optimal parameters. Note that no features need be taken from the known clean speech signals, but from the noisy speech signals only. The optimal parameters together with the corresponding features can be treated as a tuple for each relevant frame, and all of the tuples can be sequentially written to a single file.

[0040] Applying the parametric estimation process on a testing data set yielded better estimation results than the codebook approach which is also much more expensive in terms of memory and computation power requirements. Even a limited amount of available training data sufficed to improve the results. While the “average” estimator gave a good reference even for some of the badly estimated frequency bins, it failed to provide a good estimation under low-SNR conditions. After un-biasing, it was still the best choice in situations with high SNR due to its very low complexity. The codebook approach spread the mean error evenly over the spectrum and thus performed better in very low frequencies where other estimators have difficulties, especially the non-parametric ones. It did not profit very much from un-biasing and had a singular advantage at 100 km/h in very low frequencies. It was not obvious why this was the case. Its demanding complexity in memory and computation power was not justified by its performance. Under high-noise

conditions, the “parabola” and “line” estimators yielded the best results of the tested estimators, the latter slightly outperforming the former. However, the “parabola” estimation profited more from un-biasing than the “line” method. This makes it a viable choice, since the absolute computation power needed is only slightly higher than that of the “line” estimator.

[0041] Embodiments of the invention may be implemented in whole or in part in any conventional computer programming language such as VHDL, SystemC, Verilog, ASM, etc. Alternative embodiments of the invention may be implemented as pre-programmed hardware elements, other related components, or as a combination of hardware and software components.

[0042] Embodiments can be implemented in whole or in part as a computer program product for use with a computer system. Such implementation may include a series of computer instructions fixed either on a tangible medium, such as a computer readable medium (*e.g.*, a diskette, CD-ROM, ROM, or fixed disk) or transmittable to a computer system, via a modem or other interface device, such as a communications adapter connected to a network over a medium. The medium may be either a tangible medium (*e.g.*, optical or analog communications lines) or a medium implemented with wireless techniques (*e.g.*, microwave, infrared or other transmission techniques). The series of computer instructions embodies all or part of the functionality previously described herein with respect to the system. Those skilled in the art should appreciate that such computer instructions can be written in a number of programming languages for use with many computer architectures or operating systems. Furthermore, such instructions may be stored in any memory device, such as semiconductor, magnetic, optical or other memory devices, and may be transmitted using any communications technology, such as optical, infrared, microwave, or other transmission technologies. It is expected that such a computer program product may be distributed as a removable medium with accompanying printed or electronic documentation (*e.g.*, shrink wrapped software), preloaded with a computer system (*e.g.*, on system ROM or fixed disk), or distributed from a server or electronic bulletin board over the network (*e.g.*, the Internet or World Wide Web). Of course, some embodiments of the invention may be implemented as a combination of both

software (*e.g.*, a computer program product) and hardware. Still other embodiments of the invention are implemented as entirely hardware, or entirely software (*e.g.*, a computer program product).

[0043] Although various exemplary embodiments of the invention have been disclosed, it should be apparent to those skilled in the art that various changes and modifications can be made which will achieve some of the advantages of the invention without departing from the true scope of the invention.

CLAIMS

What is claimed is:

1. A computer-implemented method employing at least one hardware implemented computer processor for speech signal processing, the method comprising:
 - transforming an input microphone signal into corresponding speech component signals;
 - applying noise reduction to the speech component signals to generate noise reduced speech component signals;
 - producing corresponding synthesized speech component signals for distorted speech component signals; and
 - adaptively combining the noise reduced speech component signals and the synthesized speech component signals based on signal to noise conditions and extracted features from the input to generate an enhanced speech component signals for automatic speech recognition and hands free speech communication.
2. The method according to claim 1, wherein the speech component signals are time domain speech component signals.
3. The method according to claim 1, wherein the speech component signals are frequency domain speech component signals.
4. The method according to claim 1, wherein a non-linear function of speech components is used for producing the synthesized speech component signals.
5. The method according to claim 4, wherein the non-linear function operates to restore distorted or missing harmonics of the distorted speech component signals.
6. The method according to claim 1, wherein relevant features extracted from the distorted speech component signals are used for producing the synthesized speech component signals.
7. The method according to claim 1, wherein a source-filter model is used for producing the synthesized speech component signals.

8. The method according to claim 1, wherein a parametric model of envelope shape is used for producing the synthesized speech component signals.
9. The method according to claim 1, wherein producing corresponding synthesized speech component signals includes noise suppression pre-processing of the distorted speech component signals to provide speech synthesis reference signals.
10. The method according to claim 1, wherein producing corresponding synthesized speech component signals includes dividing the distorted speech component signals into different frequency bands.
11. A speech signal processing system employing at least one hardware implemented computer processor, the system comprising:
- a signal pre-processor for transforming an input microphone signal into corresponding speech component signals;
 - a noise suppression module for applying noise reduction to the speech component signals to generate noise reduced speech component signals;
 - a speech reconstruction module for producing corresponding synthesized speech component signals for distorted speech component signals; and
 - a signal combination block for adaptively combining the noise reduced speech component signals and the synthesized speech component signals based on signal to noise conditions to generate enhanced speech component signals for automatic speech recognition and hands free speech communication.
12. The system according to claim 11, wherein the speech component signals are time domain speech component signals.
13. The system according to claim 11, wherein the speech component signals are frequency domain speech component signals.
14. The system according to claim 11, wherein the speech reconstruction module uses a non-linear function of speech components for producing the synthesized speech component signals.

15. The system according to claim 14, wherein the non-linear function operates to restore distorted or missing harmonics of the distorted speech component signals.
16. The system according to claim 11, wherein the speech reconstruction module uses relevant features extracted from the distorted speech component signals for producing the synthesized speech component signals.
17. The system according to claim 11, wherein the speech reconstruction module uses a source-filter model for producing the synthesized speech component signals.
18. The system according to claim 11, wherein the speech reconstruction module uses a parametric model of envelope shape for producing the synthesized speech component signals.
19. The system according to claim 11, wherein the speech reconstruction module includes a noise pre-processing module for noise suppression of the distorted speech component signals to provide speech synthesis reference signals.
20. The system according to claim 11, wherein the speech reconstruction module divides the distorted speech component signals into different frequency bands.





FIG. 1

2/3

**FIG. 2**




FIG. 3A




FIG. 3B