

(51) International Patent Classification:
G10L 11/04 (2006.01)(21) International Application Number:
PCT/US2010/047888(22) International Filing Date:
3 September 2010 (03.09.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/240,062 4 September 2009 (04.09.2009) US(71) Applicant (for all designated States except US): MAS-
SACHUSETTS INSTITUTE OF TECHNOLOGY
[US/US]; 77 Massachusetts Avenue, Cambridge, MA
02139 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): WANG, Tianyu
[US/US]; 2993 Tynewick Walk, Roswell, GA 30075
(US). QUATIERI, Thomas, F., Jr. [US/US]; 15 Adams
Street, Newtonville, MA 02460 (US).(74) Agents: SMITH, James, M. et al.; Hamilton, Brook,
Smith & Reynolds, P.C., 530 Virginia Rd, P.O. Box
9133, Concord, MA 01742-9133 (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, 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, PE, PG, PH, PL, PT, RO, RS, RU, 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, SD, SL, SZ, TZ, UG,
ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, 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: METHOD AND APPARATUS FOR AUDIO SOURCE SEPARATION

FIG. 1C(a)




FIG. 1C(b)




FIG. 1C(c)





FIG. 1C(d)



(57) **Abstract:** The present invention relates to co-channel audio source separation. In one embodiment a first frequency-related representation of plural regions of the acoustic signal is prepared over time, and a two-dimensional transform of plural two-dimensional localized regions of the first frequency-related representation, each less than an entire frequency range of the first frequency related representation, is obtained to provide a two-dimensional compressed frequency-related representation with respect to each two dimensional localized region. For each of the plural regions, at least one pitch is identified. The pitch from the plural regions is processed to provide multiple pitch estimates over time. In another embodiment, a mixed acoustic signal is processed by localizing multiple time-frequency regions of a spectrogram of the mixed acoustic signal to obtain one or more acoustic properties. A separate pitch estimate of each of the multiple acoustic signals at a time point are provided by combining the one or more acoustic properties. At least one of the multiple acoustic signals is recovered using the separate pitch estimates.

- 1 -

METHOD AND APPARATUS FOR AUDIO SOURCE SEPARATION

RELATED APPLICATION

This application claims the benefit of U.S. Provisional Application No. 61/240,062, filed on September 4, 2009. The entire teachings of the above
5 application are incorporated herein by reference.

GOVERNMENT SUPPORT

The invention was supported, in whole or in part, by a grant FA 8721-05-C-0002 from the United States Air Force. The Government has certain rights in the invention.

10 BACKGROUND

Co-channel audio source separation is a challenging task in audio processing. For audio sources exhibiting acoustic properties, such as pitch, current methods operate on short-time frames of mixture signals (e.g., harmonic suppression, sinusoidal analysis, modulation spectrum [1-3]) or on single units of a time-
15 frequency distribution (e.g., binary masking [4]).

Estimating the pitch values of concurrent sounds from a single recording is a fundamental challenge in audio processing. Typical approaches involve processing of short-time and band-pass signal components along single time or frequency dimensions [12].

20 The Grating Compression Transform (GCT) has been explored [5-8] primarily for single-source analysis and is consistent with physiological modeling studies implicating 2-D analysis of sounds by auditory cortex neurons [9]. Ezzat *et al.* performed analysis and synthesis of a single speaker as the source using two-dimensional (2-D) demodulation of the spectrogram [7]. In [8], an alternative 2-D
25 modulation model for formant analysis was proposed. Phenomenological

observations in [5, 6] also suggest that the GCT invokes *separability* of multiple sources.

In [10], the GCT's ability in analysis of multi-pitch signals is demonstrated. Finally, U.S. Patent No.: 7,574,352 to Thomas F. Quatieri, Jr., the teachings of which are incorporated by reference in its entirety, relates to determining pitch estimates of voiced speech [13].

Fig. 1A illustrates schematic diagrams of 2-D frequency-related representation and compressed frequency-related representations as described in [13]. Fig. 1A demonstrates an acoustic signal waveform 104. In order to obtain pitch estimates, a first frequency-related representation of the acoustic signal is obtained over time 106. A 2-D compressed frequency-related representation 108 is then obtained by performing a 2-D transform 107 on the first frequency-related representation 106 with respect to a localized portion 110. The localized portions 110 are illustrated as windows. Each window 110 includes a frequency range that is substantially less than the frequency range of the first representation. The compressed frequency related representations 108 that result from taking a two dimensional transform of the localized portions 110 correspond to respective windows 110 in the frequency-related representation of the original waveform (shown by arrows connecting each compressed frequency-related representation to its corresponding window).

Fig. 1B is a flow diagram of the process of speaker separation described in [13]. A frequency-related representation 105 of an acoustic signal 104 is determined. This may be done in a spectrogram. The spectrogram provides a first frequency-related representation of the acoustic signal is obtained over time. Subsequently, 2-D transform 107, such as a 2-D GCT transform, is performed on the first frequency-related representation 106 with respect to localized portions. The resulting signal 108 is processed 109 to produce a speaker separated speech signal 129. Specifically, the GCT is analyzed to find a location with a maximum value and the distance from the GCT origin to the maximum value is determined 125. The reciprocal of the distance 126 is computed to produce a pitch estimate 127. The pitch estimate is then used 208 to produce a speaker separated speech signal 129.

SUMMARY

Certain example embodiments of the present invention relate to processing an acoustic signal using a first frequency-related representation of an acoustic signal prepared over time and computing a two-dimensional transform of plural two-dimensional localized regions of the first frequency-related representation, each less than an entire frequency range of the first frequency related representation, to provide a two-dimensional compressed frequency-related representation with respect to each two dimensional localized region. At least one pitch for each of the plural regions is identified. The identified pitch from the plural regions is processed to provide multiple pitch estimates over time.

On the same or alternative embodiments, a mixed acoustic signal including multiple acoustic signals may be processed. Multiple time-frequency regions of a spectrogram of the mixed acoustic signal are localized to obtain pitch candidates and provide separate pitch estimates of each of the multiple acoustic signals as a function of combining the pitch candidates. The multiple time-frequency regions may be of predetermined fixed or variable sizes. At least one of the multiple acoustic signals is recovered as a function of the separate pitch estimate.

The acoustic signal may be any audio source or a mixture of sources. For example, the acoustic signal may be a pitch-based audio source or a mixture of sources. The acoustic signal may be a speech signal. The speech signal may include a plurality of speech signals from independent speech signal sources. The two-dimensional transform may be a two-dimensional Fourier Transform.

The example embodiments may identify acoustic properties corresponding to the at least one pitch for each of the plural regions and provide the multiple pitch estimates as a function of processing the acoustic properties (e.g., pitch and the pitch dynamics). The at least one pitch may be represented as a function of a vertical distance of a representation of the at least one pitch from an origin of a frequency-related region. The frequency related region may be a Grating Compression Transform region. The at least one pitch may be represented as a function of a vertical distance and a radial angle of a representation of the at least one pitch from

an origin of a frequency-related region. The near DC components of the two-dimensional compressed frequency-related representation may be removed.

The example embodiments may identify at least one pitch for each localized time-frequency region of the spectrogram. Individual acoustic signal contents may
5 be demodulated using pitch information to recover information corresponding to an individual acoustic signal. The recovered information of the localized regions may be combined and used to reconstruct the at least one of the multiple acoustic signals.

A sinusoidal demodulation scheme may be used to demodulate individual speaker contents.

10 Certain example embodiments may be employed in processing and source separation of an acoustic signal. The acoustic signal may include multiple signals from a variety of independent sources. The acoustic signal may include multiple audio signals, multiple sounds, multiple speech signals, a mixture of speech and unvoiced acoustic signals, voiced and/or unvoiced signals combined with noise,
15 multiple unvoiced speech signals, and etc.

Certain embodiments of the present invention relate to processing a mixed acoustic signal including multiple signals. The example embodiments localize multiple time-frequency regions of a spectrogram of the mixed acoustic signal to obtain one or more acoustic properties of the mixed signal and provide a separate
20 pitch estimate of each of the multiple signals as a function of combining the one or more acoustic properties. At least one of the multiple acoustic signals is recovered based on the separate pitch estimate.

The multiple signals may include two or more unvoiced signals, two or more voiced signals, one or more unvoiced signal and a noise signal, and/or one or more
25 voiced signal and a noise signal.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing will be apparent from the following more particular description of example embodiments of the invention, as illustrated in the accompanying drawings in which like reference characters refer to the same parts
30 throughout the different views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating embodiments of the present invention.

Figs. 1A-1B include plots that demonstrate localized time-frequency regions.

Fig. 1C(a) illustrates schematic of full STFTM with localized time-frequency region centered at t_{center} and f_{center} for GCT analysis.

Fig. 1C(b) illustrates localized region of Fig. 1C(a) with harmonic structure and envelope.

Fig. 1C(c) illustrates GCT of the schematic shown in Fig. 1C(a) with baseband and modulated versions of the envelope.

Fig. 1C(d) illustrates demodulation performed to recover near-DC terms.

Fig. 2a illustrates localized time-frequency region of STFTM computed on a mixture of two speakers including a speaker with rising pitch and falling formant structure and a speaker with stationary pitch and stationary formant.

Fig. 2b illustrates GCT of the schematic shown in Fig. 2(a) including overlap of near-DC terms with speakers exhibiting the same vertical distances from the ω -axis corresponding to equal pitch values and separability being maintained due to distinct angular positions off of the Ω -axis.

Fig. 2c illustrates Demodulation to recover near-DC terms of one speaker.

Figs. 3a-3d are plots that demonstrate results obtained using analysis and synthesis of a single speaker.

Figs. 4a-4d are plots that demonstrate results obtained using analysis and synthesis of two speakers.

Fig. 5 is a table that includes average signal-to-noise ratios (SNR) of the original and reconstructed waveforms shown in Figs. 3 and 4.

Fig. 6a is a plot of an impulse train with linearly increasing pitch that illustrates a spectrogram with localized regions across frequency for a single time segment.

Fig. 6b demonstrates a zoomed-in region from Fig. 6(a) with vertical $(2\pi/\omega_s \cos\theta)$ and spatial $(2\pi/\omega_s)$ distance between harmonic lines and 2-D sine wave orientation (θ) .

Fig. 6c demonstrates magnitude of GCT of the schematic shown in Fig. 6b with $\omega_s, \theta, \omega_s \cos\theta$. The DC component is removed for display purposes.

Fig. 6d illustrates pitch estimates obtained from vertical vs. radial distances to GCT origin.

Fig. 7a includes a plot that relate to a localized region with two distinct pitch values with no temporal change.

5 Fig. 7b illustrates GCT corresponding to the schematic shown in Fig. 7a with separability occurring along the Ω -axis.

Fig. 7c illustrates localized region with two pitch candidates with the *same* pitch value but different temporal dynamics.

10 Fig. 7d illustrates GCT corresponding to the schematic shown in Fig. 7c, with separability occurring from difference in temporal dynamics.

Fig. 8a demonstrates examples of analysis of two synthesized concurrent vowels with formant structure of rising and falling-pitched vowels in Region 1 (R1) and Region 2 (R2).

15 Fig. 8b illustrates $r_{xx}[n]$ for R1, shown in Fig. 8a, with lag corresponding to dominant pitch.

Fig. 8c illustrates $r_{xx}[n]$ for R2 shown in Fig. 8a.

Fig. 8d illustrates spectrogram of concurrent vowels.

Fig. 8e illustrates GCT for R1 (shown in Fig. 8a) with dominant pitch.

Fig. 8f illustrates GCT for R2 with two sets of pitch peaks.

20 Fig. 9 illustrates post-processing methods for assigning pitch value at time *. The term “s” denotes single low-frequency region, dashed regions denote regions used in clustering for synthetic speech, and shaded regions denote regions used in clustering for real speech.

Fig. 10 demonstrate estimates for the synthetic mixtures including concurrent 25 vowel pitch estimates, *clustering*, *single*, and true pitch values. Fig. 10a illustrates estimates rising 125-175 Hz + constant 150 Hz. Fig. 10b illustrates estimates falling 250-200 Hz + rising 125-150 Hz. Fig. 10c illustrates estimates rising 100-150 Hz + rising 175-225 Hz. Fig. 10d illustrates estimates falling 200-150 Hz + rising 150-200 Hz. Estimates start (end) at 50 (250) ms to remove edge effects.

30 Fig. 11 demonstrates estimates for real speech mixture with no crossing pitch tracks. In Fig. 11a, all-voiced mixture spectrogram with separate pitch trajectories

are illustrated. The all voice mixture includes Male – “Why were you away a year?” + Female – “Nanny may know my meaning.” The first 250 ms and last 50 ms excluded to remove edge effects in clustering due to initial and final silent regions. In Figs. 11b-11d, *truth*, *oracle*, *clustering*, and *single* estimates are illustrated.

5 Fig. 12a demonstrates estimates for real speech mixture with crossing pitch tracks. All-voiced mixture spectrogram with separate pitch trajectories, Male - "Why were you away a year?" + Female - "Nanny may know my meaning." The first 250 ms and last 50 ms are excluded to remove edge effects in clustering due to initial and final silent regions. Figs. 12b-12d illustrate estimates of *truth*, *oracle*,
10 *clustering*, *single* . In Fig. 12c a “jump” due to differences in energy between sources in localized time-frequency regions is illustrated.

Fig. 13 is a table that includes average errors for examples studied in Figs. 10-12.

Fig. 14 includes plot of candidate collections obtained across frequency
15 regions from GCT analysis. Each time-frequency region's candidate is represented by * and reference pitch and pitch-derivative values are shown using solid lines.

Figs. 15a-15d includes schematics that illustrate a method for obtaining pitch candidates according to an example embodiment.

Fig. 16(a) illustrates a collection of histogram slices of pitch candidates and
20 Fig. 16(b) illustrates a single histogram slice.

Fig. 17 demonstrates the GCT computed for a region in which the formant structure of one speaker in the mixture exhibits a formant peak. Fig. 17a demonstrates analysis of spectrogram of mixture and localized region. Fig. 17b is a zoomed-in version of localized region shown in Fig. 17a. Fig. 17c demonstrates
25 GCT computed without minimizing near-DC components. Fig. 17d demonstrates GCT computed with near-DC components minimized.

Fig. 18 demonstrates the GCT computed for a region in which the formant structures of both speakers in the mixture exhibit formant valleys. Fig. 18a illustrates spectrogram of analyzed mixture and localized region. Fig. 18b
30 demonstrates a zoomed-in version of the localized region of Fig. 18a. Fig. 18c demonstrates GCT computed without minimizing near-DC components. Fig. 18d demonstrates GCT computed with near-DC components minimized..

Fig. 19 includes results of LDA analysis of features on training data and maximum-likelihood fits to Gaussian distributions.

Fig. 20 includes plots of performance metrics obtained using a hypothesis testing framework on testing data.

5 Figs. 21a-21d include histogram plots of pitch candidates.

Figs. 22a-22b include examples of pruned (Fig. 22a) and raw (Fig. 22b) pitch candidates and resulting median-based clustering method to assign pitch candidates to time points.

10 Fig. 23a-23b illustrate examples of pruned (Fig. 23a) and raw (Fig. 23b) pitch candidates and resulting k-means clustering method to assign pitch candidates to time points.

Fig. 24 includes a table of average metrics across all test mixtures for a median clustering method.

15 Fig. 25 includes a table of average metrics across all test mixtures for a k-means clustering method.

Fig. 26A is an illustration of a multi-pitch estimation method that employs clustering and Kalman filtering.

20 Fig. 26B is a plot of log of narrowband short-time Fourier Transform magnitude of mixture of "Walla Walla" and "Lawyer" sentences spoken by a male and female speaker.

Fig. 26C is an illustration of band-wise classification performance of linear discriminate analysis on test data.

Fig. 26D is an illustration of resulting binary mask of pruning of the plot shown in Fig. 27A with 1's and 0's.

25 Fig. 26E is a plot that illustrates the average root mean square and standard errors for data collected from eight males and eight females and processed using embodiments of the present invention that employ Kalman filtering and clustering for multi-pitch analysis.

30 Figs. 26F-26K are plots that illustrate estimation results obtained from embodiments of the present invention that employ Kalman filtering and clustering for multi-pitch analysis.

Fig. 26L illustrates a narrowband spectrogram having a local region.

Fig. 26 M illustrates a zoomed-in portion of the local regions shown in Fig. 26L.

Fig. 26N illustrates the GCT representation with near-DC component to be removed for analysis and synthesis and series modulated components.

5 Fig. 26O illustrates demodulation in the GCT domain for analysis and synthesis using series to reconstruct near-DC terms.

Figs. 27A-27D are four spectrograms of synthetic signals.

Fig. 28 is a table that outlines the SNR value for all four signals shown in Fig. 27A-27D across four distinct regions sizes.

10 Fig. 29A illustrates the result of base tiling showing base region and its neighbors.

Fig. 29B illustrates schematic of base region grown from (a) and its new neighbors.

15 Figs. 30A and 30B are high-level illustrations of an audio separation system according to an example embodiment.

Figs. 31A and 31B are high-level illustrations of a two-dimensional multi-pitch estimation subsystem according to an example embodiment.

Figs. 32A and 32B are high-level illustrations of a two-dimensional source separation subsystem according to an example embodiment.

20 Fig. 33 is a high-level illustration of a system for processing an acoustic signal according to example embodiments of the present invention.

Fig. 34 is a high-level illustration of a system for processing a mixed acoustic signal according to example embodiments of the present invention.

DETAILED DESCRIPTION

25 The teachings of all patents, published applications and references cited herein are incorporated by reference in their entirety.

Certain example embodiments of the present invention address multi-pitch estimation and speaker separation using a two-dimensional (2-D) processing framework of a mixture signal. For the speaker separation task, one example
30 embodiment of the present invention relates to a method and corresponding apparatus for two-dimensional (2-D) processing approach for co-channel speaker

- separation of voiced speech. Localized time-frequency regions of a narrowband spectrogram may be analyzed using 2-D Fourier transforms to determine a 2-D amplitude modulation model based on pitch information for single and multi-speaker content in each region. Harmonically-related speech content may be mapped to concentrated entities in a transformed 2-D space, thereby motivating 2-D demodulation of the spectrogram for analysis/synthesis and speaker separation. Using a priori pitch estimates of individual speakers, the example embodiment determines through a quantitative evaluation: 1) utility of the model for representing speech content of a single speaker and 2) its feasibility for speaker separation.
- Localized time-frequency regions of a narrowband spectrogram may be analyzed using 2-D Fourier transforms. This representation is also referred to as the Grating Compression Transform (GCT).

Speaker Separation Using A Priori Pitch Estimates of Individual Audio Sources

Single-speaker Modeling

- Fig. 1C includes plots that demonstrate localized time-frequency regions according to certain example embodiments. A localized time-frequency region $s[n, m]$ (discrete-time and frequency n, m) of a narrowband short-time Fourier transform magnitude (STFTM) for a single voiced utterance may be computed. Fig. 1C(a) illustrates a schematic 101 of full STFTM with localized time-frequency region centered at t_{center} 103 and f_{center} 102 for GCT analysis (area shown using rectangle 110). The localized region 110 is shown in more details in the schematic of Fig. 1b. As demonstrated in Fig. 1C(b), the localized region 110 includes a harmonic structure (parallel lines) and an envelope (shaded area 120). The localized region 110 includes a frequency range that is substantially less than the frequency range of the first representation. A 2-D amplitude modulation (AM) model from [8] may be extended such that:

$$\begin{aligned} s[n, m] &\approx (\alpha_0 + \cos(\Phi[n, m]))a[n, m] \\ \Phi[n, m] &= \omega_s(n \cos \theta + m \sin \theta) + \varphi. \end{aligned} \quad (1)$$

Thus, a sinusoid with spatial frequency ω_s , orientation θ , and phase φ rests on a DC pedestal α_0 and modulates a slowly-varying envelope $a[n, m]$. The 2-D Fourier transform of $s[n, m]$ (i.e., the GCT) is

$$S(\omega, \Omega) = \alpha_0 A(\omega, \Omega) + 0.5e^{-j\varphi} A(\omega + \omega_s \sin \theta, \Omega - \omega_s \cos \theta) + 0.5e^{j\varphi} A(\omega - \omega_s \sin \theta, \Omega + \omega_s \cos \theta) \quad (2)$$

5 where ω and Ω map to n and m , respectively. The sinusoid represents the harmonic structure associated with the speaker's pitch [5, 10]. Denoting f_s as the waveform sampling frequency and N_{STFT} as the discrete-Fourier transform (DFT) length of the STFT, the GCT parameters relate to the speaker's pitch (f_0) at the center (in time) of $s[n, m]$ (as shown in Figs. 1C(b) and 1C(c)) [5, 10]:

$$10 \quad f_0 = (2\pi f_s) / (N_{STFT} \omega_s \cos \theta). \quad (3)$$

A change in f_0 (Δf_0) across Δn results in an *absolute* change in frequency of the k^{th} pitch harmonic by $k\Delta f_0$. Therefore, in a localized time-frequency region (Fig. 1C(b))

$$\tan \theta \approx (k\Delta f_0) / \Delta n. \quad (4)$$

For a particular $s[n, m]$ with center frequency f_{center} (Figure 1a), f_0 can be obtained
15 from (3) such that $k \approx f_{\text{center}} / f_0$. The rate of change of f_0 ($\partial f_0 / \partial t$) in $s[n, m]$ is then

$$\partial f_0 / \partial t \approx \Delta f_0 / \Delta n = (f_0 \tan \theta) / f_{\text{center}}. \quad (5)$$

Finally, φ corresponds to the position of the sinusoid in $s[n, m]$; for a non-negative DC value of $a[n, m]$, φ can be obtained by analyzing the GCT at $(\omega = \omega_s \sin \theta, \Omega = \omega_s \cos \theta)$
:

$$20 \quad \varphi = \text{angle}[S(\omega_s \sin \theta, \omega_s \cos \theta)]. \quad (6)$$

Fig. 1C(c) demonstrates concentrated entities 121 in the GCT near DC and at 2-D carriers. As shown in Fig. 1C(c), harmonically related speech content in each $s[n, m]$ are mapped into *concentrated* entities 121 in the GCT near DC and at 2-D “carriers.”

25 Fig. 1C(d) illustrates demodulation process 122 performed to recover near-DC terms. Once the near-DC terms are removed or corrupted, approximate *recovery* of the near-DC terms from the carrier terms using sinusoidal demodulation becomes

possible. Using demodulation, the full STFTM may then be recovered and combined with the STFT phase for approximate waveform reconstruction.

Multi-speaker Modeling

- 5 Certain example embodiments approximate the STFTM computed for a mixture of N speakers in a localized time-frequency region $x[n, m]$ as the sum of their individual magnitudes. Using the model of (1),

$$x[n, m] \approx \sum_{i=1}^N \alpha_{0,i} a_i[n, m] + \sum_{i=1}^N a_i[n, m] \cos(\omega_i [n \cos \theta_i + m \sin \theta_i] + \varphi_i). \quad (7)$$

- Equation (7) invokes the sparsity of harmonic line structure from distinct speakers in the STFTM (i.e., when harmonic components of speakers' are located at different frequencies). Nonetheless, separation of speaker content in the GCT may still be maintained when speakers exhibit harmonics located at identical frequencies (e.g., due to having the same pitch values, when pitch values are integer multiples of each other) due to its representation of pitch dynamics through θ in (7) [10].

- 15 Fig. 2 illustrates an example embodiment that employs two speakers with equal pitch values and distinct pitch dynamics. In Fig. 2a, a localized time-frequency region 201 of STFTM is computed on a mixture of two speakers. One speaker may include a rising pitch and a falling formant structure 220. The other speaker may include a stationary pitch and a stationary formant structure 210. The speakers may give the same pitch value at center of region (shown by arrow 230). The GCT 202 of the localized time-frequency region 201 is illustrated in Fig. 2. The GCT 202 demonstrates overlap of near-DC terms 240. The two speakers exhibit the same vertical distances 250 from the ω -axis corresponding to equal pitch values. However, separability is maintained due to distinct angular positions off of the Ω -axis. Thus, as shown in Figs. 2a-2b, since the two speakers have equal pitch values but distinct pitch dynamics, separability in the GCT is possible.

The 2-D Fourier transform of (7) is

$$X(\omega, \Omega) = \sum_{i=1}^N \alpha_{0,i} A_i(\omega, \Omega) + 0.5 \sum_{i=1}^N A_i(\omega + \omega_i \sin \theta_i, \Omega - \omega_i \cos \theta_i) e^{-j\theta_i} + 0.5 \sum_{i=1}^N A_i(\omega - \omega_i \sin \theta_i, \Omega + \omega_i \cos \theta_i) e^{j\theta_i}. \quad (8)$$

For slowly-varying $A_i(\omega, \Omega)$, the contribution to $X(\omega, \Omega)$ from multiple speakers exhibits *overlap* near the GCT origin (Figure 2b); however, as in the single-speaker case, $A_i(\omega, \Omega)$ can be estimated through sinusoidal demodulation according to the proposed model. This model therefore motivates localized 2-D demodulation of the STFTM computed for a mixture of speakers for the speaker separation task (Figure 2c). By demodulating the STFTM near-DC terms of a speaker may be recovered.

In certain embodiments, a sinusoidal demodulation in conjunction with a least-squared error fit may be employed to estimate the gains $\alpha_{0,i}$ in (7) and (8). This notion can be generalized to include any parametric representation and can be extended beyond the gains $\alpha_{0,i}$ to include a parametric representation of the entire amplitude $a[n, m]$ and phase $\varphi[n, m]$ functions of each signal component in the GCT domain.

15 Single-speaker Analysis and Synthesis

To assess the AM model's ability to represent speech content of a single speaker, an STFT is computed for the signal using a 20 milliseconds (ms) Hamming window, 1-ms frame interval, and 512-point DFT. From the full STFTM ($s_F[n, m]$), localized regions centered at k and l in time and frequency ($s_{kl}[n, m]$) of size 625 Hz by 100 ms are extracted using a 2-D Hamming window ($w_h[n, m]$) for GCT analysis. A high-pass filter $h_{hp}[n, m]$ is applied to each $s_{kl}[n, m]$ to remove $\alpha_0 A(\omega, \Omega)$ in (2) (this result is denoted as $s_{kl, hp}[n, m]$). The $h_{hp}[n, m]$ is a circular filter with cut-offs at $\omega = \Omega = 0.1\pi$, corresponding in ω to a ~ 300 Hz upper limit of f_0 values observed in analysis.

25 For each $s_{kl, hp}[n, m]$, certain example embodiments may approximately recover $\alpha_0 A(\omega, \Omega)$ using 2-D sinusoidal demodulation. The carrier ($\cos(\Phi[n, m])$) parameters are determined from the speaker's pitch track using (3) for ω_s and (6) for φ . To determine θ , a linear least-squared error fit is applied to the pitch values spanning the 100-ms duration of $s_{kl, hp}[n, m]$. The slope of this fit approximates $\partial f_0 / \partial t$ such that

θ is estimated using (5). The term $s_{kl, hp}[n, m]$ is multiplied by the carrier generated from these parameters followed by filtering with a circular low-pass filter $h_{lp}[n, m]$ with cut-offs at $\omega = \Omega = 0.1\pi$ (this result is denoted as $\hat{a}[n, m]$). The term $\hat{a}[n, m]$ is combined with the carrier using (1) and set equal to $s_{kl}[n, m]$

$$s_{kl}[n, m] = (\alpha_0 + \cos(\Phi[n, m]))\hat{a}[n, m]. \quad (9)$$

For each time-frequency unit of $s_{kl}[n, m]$, (9) corresponds to a linear equation in α_0 since the values of $s_{kl}[n, m]$, $\hat{a}[n, m]$, and $\cos(\Phi[n, m])$ are known. This over-determined set of equations is solved in the least-squared error (LSE) sense. The resulting estimate of $s_{kl}[n, m]$ using the estimated α_0 , $\hat{a}[n, m]$, and $\cos(\Phi[n, m])$ is denoted as $\hat{s}_{kl}[n, m]$. The *full* STFTM estimate $\hat{s}_F[n, m]$ is obtained using overlap-add (OLA) with a LSE criterion (OLA-LSE) [11]

$$\hat{s}_F[n, m] = \frac{\sum_k \sum_l w_h[kT - n, lF - m] \hat{s}_{kl}[n, m]}{\sum_k \sum_l w_h^2[kT - n, lF - m]}. \quad (10)$$

OLA step sizes in time and frequency (T and F) are set to $1/4$ of the size of $w_h[n, m]$. The term $\hat{s}_F[n, m]$ is combined with the STFT phase for waveform reconstruction using OLA-LSE [11].

Speaker Separation

For speaker separation, the demodulation steps are nearly identical to those used for single speaker analysis and synthesis. The demodulation steps are applied to the *mixture* signal. Assuming that $x_{kl}[n, m]$ is a localized region of the full STFTM computed for the *mixture* signal centered at k and l in time and frequency, the term $x_{kl}[n, m]$ is filtered with $h_{hp}[n, m]$ to remove the overlapping $\alpha_{0,i} A_i(\omega, \Omega)$ terms at the GCT origin (this result is denoted as $x_{kl, hp}[n, m]$). A cosine carrier for each speaker is generated using the corresponding pitch track and multiplied by $x_{kl, hp}[n, m]$ to obtain

$$\begin{aligned} x_{kl,i}[n, m] &= x_{kl, hp}[n, m] \cos(\omega_i [n \sin \theta_i + m \cos \theta_i] + \varphi_i) \\ &= \hat{a}_i[n, m] + c[n, m]. \end{aligned} \quad (11)$$

If the speakers' carriers are in distinct locations of the GCT, $c[n, m]$ summarizes cross terms away from the GCT origin such that $\hat{a}_i[n, m]$ can be obtained

by filtering $x_{kl,i}[n,m]$ with $h_p[n,m]$. For each speaker, $\hat{a}_i[n,m]$ is combined with its respective carrier using (1). These results are summed and set equal to $x_{kl}[n,m]$ to solve for $\alpha_{0,i}$ in the LSE sense:

$$x_{kl}[n,m] = \sum_{i=1}^N (\alpha_{0,i} + \cos(\Phi[n,m])) \hat{a}_i[n,m] \quad (12)$$

- 5 Since GCT represents pitch *and* pitch dynamics; it may, therefore, invoke improved speaker separability over representations relying solely on harmonic sparsity. In a region where speakers have equal pitch values and the same temporal dynamics, however, (12) invokes a near-singular matrix. To address this, the angle between the $\hat{a}_i[n,m]$ columns of the matrix may be computed. When this angle is below a
- 10 threshold (in certain example embodiments this threshold may be $\pi/10$), the $\alpha_{0,i}$ is solved for by reducing the matrix rank to that corresponding to a single speaker.

Finally, the estimated full STFTMs of the target speakers are reconstructed using (10). Speaker waveforms are then reconstructed using OLA-LSE by combining the estimated STFTMs with the STFT phase of the mixture signal.

- 15 In certain embodiments, the acoustic signal being processed may include multiple voiced signals and unvoiced signals. In such embodiments, only the voice signal components require pitch estimation.

- Certain embodiments may modify an audio-signal component in the GCT space prior to reconstruction (e.g., pitch modification to transfer/modify a
- 20 concentrated entity).

Evaluation

- Two all-voiced sentences sampled at 8 kHz (“Why were you away a year, Roy?” and “Nanny may know my meaning”), spoken by 10 males and females (40
- 25 total sentences), are analyzed. Pitch estimates of the individual sentences are determined prior to analysis from an autocorrelation-based pitch tracker.

First analysis and synthesis of a single speaker is performed. For comparison, a waveform by filtering $s_p[n,m]$ with an adaptive filter

$$h_s[n,m] = h_p[n,m](1 + 2 \cos(\omega_s[n \sin \theta + m \cos \theta] + \varphi_s)) \quad (13)$$

- 30 is also generated. In (13), ω_s , θ , and φ_s are determined for each localized time-

frequency region using the speaker's pitch track. The term $h_p[n, m]$ is obtained for use in single speaker analysis and synthesis. The filtered STFTM is used to recover the waveform.

To assess the feasibility of GCT-based speaker separation, mixtures of two sentences (Nanny and Roy) spoken by 10 males and females mixed at 0 dB (90 mixtures total) are analyzed. For comparison, a baseline sine wave-based separation system (SBSS) is used. The SBSS models sine wave amplitudes and phases given their frequencies (e.g., harmonics) for each speech signal [2]. This baseline is chosen for comparison as it similarly uses a priori pitch estimates to obtain the sinusoidal frequencies, and to assess potential benefits of the GCT's explicit representation of pitch dynamics.

Figs. 3a-3d are plots that demonstrate STFTMs obtained in the single-speaker experiment and a priori pitch estimates. Fig. 3a demonstrates the STFT magnitude of a single speaker sentence (Roy) 310. Fig. 3b demonstrates the recovered STFT magnitude using a control method 320. Fig. 3c demonstrates the recovered STFT magnitude using a demodulation method. Fig. 3d demonstrates the a priori pitch estimates of sentence analyzed in Fig. 3a – Fig. 3c.

In plots of Figs. 3a-3d, STFTMs is obtained in the single-speaker experiment and a priori pitch estimates. The demodulation scheme appears to provide a similar reconstruction as the control method.

Figs. 4a-4d are plots that demonstrate STFTMs obtained in the two-speaker experiment. Fig. 4a demonstrates the STFT magnitude of mixture speaker sentences (i.e., Nanny and Roy). Fig. 4b illustrates the recovered STFT magnitude of the Roy sentence using SBSS with resulting SNR listed. Fig. 4c illustrates the recovered STFT magnitude of the Roy sentence using demodulation. Fig. 4d demonstrates a priori pitch estimates of target 440 and interfering 450 speakers. The pitch tracks exhibit crossings throughout mixture.

In Fig. 4, the resulting STFTMs for the separation task using the single-speaker sentence as the target is shown. In this example, the pitch tracks of the target and interferer exhibit crossings (Fig. 4d), thereby leading to overlapping

harmonic structure in the mixture STFTM. Qualitatively, GCT demodulation appears to provide a more faithful reconstruction of the target than SBSS.

Fig. 5 is a table 500 that includes average signal-to-noise ratios (SNR) of the original and reconstructed waveforms shown in Figs. 3 and 4. In the single speaker case, demodulation provides a better reconstruction than filtering by ~1.3 dB. One possible cause for this is the introduction of negative magnitude values in the filtered STFTM. These effects are likely minimized in demodulation through the LSE fitting procedure. Nonetheless, both methods provide good reconstruction of the waveform with overall SNR > 11 dB. In the mixed speaker case, consistent with the recovered STFTMs (Fig. 4), demodulation affords a larger gain in SNR than SBSS and on average. This is due to the GCT's explicit representation of pitch dynamics. In informal listening for the single speaker case, subjects report no perceptual difference between the filtering and demodulation methods in relation to the original signal. In the mixed speaker case, subjects report intelligible reconstructions of the target speech for both methods with a reduced amplitude of the interferer. However, in assessing SBSS, subjects report that the interferer sounded "metallic" while this synthetic quality was not perceived for the GCT system. These observations demonstrate the utility of the AM model for representing speech content of a single speaker. They also demonstrate the feasibility of GCT for speaker separation and its advantages in representing pitch dynamics for this task.

Therefore, a 2-D modulation model accounting for near-DC terms of the GCT provides good representation of speech content of a single speaker and may be used for co-channel speaker separation. The present method is to be combined with the subsequently discussed method for performing multi-pitch estimation. Specifically, the *a priori* estimated pitch estimates used in this section will be replaced by those obtained using the multi-pitch estimation method.

GCT-based Multi-Pitch Estimation Methods

GCT-based Analysis of Multi-pitch Signals

For the multi-pitch estimation task, certain example embodiments may include two-dimensional (2-D) processing and analysis of multi-pitch speech sounds. Short-space 2-D Fourier transform magnitude of a narrowband spectrogram

are invoked, and harmonically-related signal components are mapped to multiple concentrated entities in a new 2-D space. First, localized time-frequency regions of the spectrogram are analyzed to extract pitch candidates. These candidates are then combined across multiple regions for obtaining separate pitch estimates of each speech-signal component at a single point in time (referred to as multi-region analysis (MRA)). By explicitly accounting for pitch dynamics within localized time segments, this separability is distinct from that which can be obtained using short-time autocorrelation methods typically employed in state-of-the-art multi-pitch tracking algorithms.

Framework

Fig. 6 includes plots that illustrate an impulse train with linearly increasing pitch. Figs. 6a-6b illustrate examples of localized time-frequency regions $s[n, m]$ (with discrete-time and frequency: n, m) of a narrowband short-time Fourier transform (STFT) log-magnitude exhibiting harmonic line structure for an impulse train with linearly increasing pitch (125-200 Hz). In Fig. 6a, spectrogram with localized regions (rectangles 610) across frequency for a single time segment (arrow 620) are shown. Fig. 6b demonstrates a zoomed-in region 610.

A 2-D sinewave model for $s[n, m]$ is [13]

$$s[n, m] \approx K + \cos(\omega_s \Phi[n, m]) \quad (14)$$

where ω_s denotes the local spatial frequency of the sinusoid, $\Phi[n, m]$ is a 2-D phase term indicating its orientation, and K is a constant DC term. The term $\Phi[n, m]$ is defined as

$$\Phi[n, m] = n \sin \theta + m \cos \theta \quad (15)$$

where θ is the angle of rotation of the harmonic lines relative to the time axis. The 2-D Fourier transform of $s[n, m]$ is then

$$S(\omega, \Omega) \approx 2\pi K \delta(\omega, \Omega) + 2\pi \delta(\omega + \omega_s \sin \theta, \Omega - \omega_s \cos \theta) + 2\pi \delta(\omega - \omega_s \sin \theta, \Omega + \omega_s \cos \theta) \quad (16)$$

such that the harmonic line structure maps to a set of impulses in the GCT (Fig. 6c). The present embodiment employs $\omega_s \cos \theta$ (vertical distance of impulses to the GCT

origin) to better represent pitch across all frequency regions ($f_{0, \text{vertical}} = \frac{1}{N_{\text{STFT}}} \frac{2\pi f_s}{\omega_s \cos \theta}$).

This is consistent with the fact that $\omega_s \cos \theta$ is inversely related to the vertical spacing between harmonic peaks in $s[n, m]$. Here, f_s is the sampling rate of the waveform, and N_{STFT} is the discrete-Fourier transform (DFT) length used to compute the spectrogram.

5 Fig. 6d demonstrates pitch estimates obtained from vertical vs. radial distances to GCT origin. Specifically, Fig. 1d compares $f_{0,radial}$ 630 and $f_{0,vertical}$ 640 computed across multiple frequency regions 610 for a local time segment 620 by peak-picking of the GCT magnitude. While $f_{0,vertical}$ 640 remains constant across frequency regions and corresponds to the true pitch value at the center of segment 10 (~ 160 Hz), $f_{0,radial}$ 630 decreases across frequency regions by ~ 10 Hz. This effect is presumably due to the increased fanning of harmonic line structure in higher-frequency regions with changing pitch. This comparison illustrates that rotation of the GCT components (i.e., θ) *increases* from low- to high-frequency regions.

15 Analysis of Multi-Pitch Signals and Separability of Pitch Information in the GCT

Extending (15) and (16) to the case of N concurrent signals,

$$s[n, m] \approx \sum_{i=1}^N (K_i + \cos(\omega_i \Phi[n, m; \theta_i])) \quad (17)$$

$$20 \quad S(\omega, \Omega) \approx 2\pi \sum_{i=1}^N K_i \delta(\omega, \Omega) + 2\pi \sum_{i=1}^N \delta(\omega + \omega_i \sin \theta_i, \Omega - \omega_i \cos \theta_i) \\ + 2\pi \sum_{i=1}^N \delta(\omega - \omega_i \sin \theta_i, \Omega + \omega_i \cos \theta_i). \quad (18)$$

Here, the (log)-magnitude STFT of a mixture of signals is approximated as the *sum* of the STFT (log)-magnitudes computed for each individual signal. This approximation holds best when the contribution to the STFT from distinct sources occupies different frequency bands. Nonetheless, separation of pitch in the GCT can 25 be maintained even when these conditions do not necessarily hold, i.e., when a frequency band contains more than one source (with or without similar pitch values).

Fig. 7 includes plots that relate to a localized region with two distinct pitch values. In Fig. 7a, a localized region with two distinct pitch values and no temporal change is illustrated. Specifically, in Fig. 7b, a region of the spectrogram 710 30 having two sets of harmonic lines corresponding to two distinct pitch trajectories

that are *constant* through $s[n, m]$ is shown. Fig. 7b illustrates the GCT plot 720 of the localized region 710 shown in Fig. 7a. The GCT exhibits two sets of impulses along the Ω -axis 702. In this case, separability can only be achieved when the two pitch values are sufficiently different. This set-up also generalizes to the case when the two trajectories in $s[n, m]$ move at the same rate and direction.

Fig. 7c illustrates that separability occurs along the Ω -axis and Fig. 7d illustrates the GCT plot of Fig. 7c. Specifically, Fig. 7c-7d illustrate a condition in which two pitch trajectories have *equal* pitch values defined at the center of $s[n, m]$ in time but are moving in opposite directions at the same rate. Despite the overlap of harmonic structure in $s[n, m]$, the GCT maintains separability of pitch information due to its *explicit* representation of the underlying temporal trajectories of the two sources in the values of θ_i (i.e., $\theta_1 = -\theta_2 = \theta$). More generally, this separability holds under conditions where the rates of change of the two pitch trajectories are different (i.e., $\theta_1 \neq \theta_2$).

Since, for moving pitch trajectories, θ increases from low- to high-frequency regions (Fig. 6d), analysis of multiple regions across frequency and time can be expected to provide better separability of pitch information than that of a *single* low-frequency region across time as in [13].

20 Comparison to Short-time Autocorrelation Analysis

The GCT's representation of pitch dynamics within a local time segment invokes separability of pitch information distinct from that obtained in short-time autocorrelation analysis.

Fig. 8 demonstrates analysis of two synthesized concurrent vowels with rising and falling pitch contours of 150-200 Hz and 200-150 Hz across a 200-ms duration. Fig. 8a demonstrates the formant structure for the vowels. In Region 1 (R1 810), the rising vowel exhibits a formant peak 812 while the falling vowel exhibits a valley 814. In Region 2 (R2 820), a formant peak 822 is present for both vowels. Analysis is done at the center of the mixture where both sources have pitch values of approximately 175 Hz.

For comparison, two linear-phase band-pass filters centered at the formant peaks of R1 810 and R2 820 were applied to the waveform. To obtain an envelope [12], filtered waveforms were then half-wave rectified and low-pass filtered (cutoff = 800 Hz). The normalized autocorrelation ($r_{xx}[n]$) was computed for a 30-ms duration of the envelopes (Figs. 8b-8c). For R1 810, a single distinct pitch estimate and its sub-harmonics are present (Fig. 8b, arrow 830). However, $r_{xx}[n]$ for R2 820 (Fig. 8c) reflects the interaction of closely-spaced periodicities and appears noisy.

Figs. 8e-8f demonstrate GCTs computed over localized time-frequency regions at R1 810 and R2 820 (Fig. 8d). A single dominant set of impulses, corresponding to a single pitch value, is present in the GCT for R1, similar to $r_{xx}[n]$ for R1 810. However, two distinct sets of peaks can be seen for R2 820 (Fig. 8e-8f, arrows 890, 895) corresponding to two similar pitch values. The GCT can, therefore, separate pitch information of two speakers with similar energies and pitch values in a localized set of frequency bands by exploiting the temporal dynamics of their underlying pitch trajectories. This separability is distinct from that obtained using short-time autocorrelation analysis (as compared with Fig. 8c). This separability may be generalized to the case where source signals exhibit similar energies but different pitch values/temporal dynamics.

20 Multi-pitch Estimation Method 1: GCT-based Multi-region Analysis Approach

The log-STFT magnitude is computed for all mixtures with a 25-ms Hamming window, 1-ms frame interval, and 512-point DFT. Time-frequency regions of size 100 ms by 700 Hz are extracted from the spectrogram at a 5-ms and 140-Hz region interval in time and frequency, respectively. A 2-D gradient operator is applied to the spectrogram prior to extraction to reduce the contribution of the DC and near-DC components to the GCT. To obtain pitch candidates for each region, the GCT magnitude is multiplied by three binary masks derived from thresholding the 1) overall amplitude, 2) gradient (∇GCT), and 3) Laplacian (ΔGCT). The thresholds are chosen as $\max(GCT)/3$, $\max(\nabla GCT)/3$, and $\min(\Delta GCT)/3$. Region growing is performed on the masked GCT, and pitch candidates are obtained by extracting the location of the maximum amplitude in each resulting region.

Candidates corresponding to the two largest amplitudes are kept for each time-frequency region. In the case where only a single pitch value is present, the value is assigned twice to the region.

5 Post-processing

For synthetic speech, a simple clustering method is used to assign pitch values at each point in time from the candidates of GCT-based MRA. All candidates at a single point in time are collected and sorted, and the median of the top and bottom halves of the collection are then chosen as the two pitch values. A similar technique is used for real speech. However, due to the longer duration of these signals, the temporal continuity of the underlying pitch contours is used in clustering. At each 5-ms interval for a time-frequency region, pitch candidates from its neighboring regions in time spanning 100 ms and across frequencies are combined for clustering. To compare GCT-based MRA with [13], each 5-ms interval is assigned the two candidates from analyzing a *single* low-frequency region.

Fig. 9 illustrates post-processing methods for assigning pitch value at time * 910. Variable “s” denotes single low-frequency region. Dashed regions 930 denote regions used in clustering for synthetic speech and shaded regions 940 denote regions used in clustering for real speech. Finally, oracle pitch values are obtained by assigning to each time point the pitch candidate from GCT-based MRA closest in frequency to the true pitch values. The accuracy of these estimates is viewed as assessing the value of GCT-based MRA for obtaining pitch candidates independent of post-processing (e.g., clustering).

25

Data Used in Evaluation of Multi-Pitch Estimation Method 1

Concurrent vowels with linear pitch trajectories spanning 300 ms are synthesized using a glottal pulse train and an all-pole formant envelope with formant frequencies of 860, 2050, and 2850 Hz and bandwidths of 56, 65, 70 Hz (/ae/) [5]. For real speech, two all-voiced sentences spoken by a male and female are used. Two cases are analyzed to illustrate typical pitch-trajectory conditions: 1) separate or 2) crossing trajectories within the utterance. All signals are mixed at 0 dB overall

30

signal-to-signal ratio (SSR) and pre-emphasized prior to analysis. True pitch value are obtained using a single-pitch estimator on the signals prior to mixing [6].

Results

5 Figs. 10-12 demonstrate estimates for the synthetic and real speech mixtures. The total-best percent error between estimates and truth for both source signals is computed at each time point:

$$\%error = 100 \left[\frac{|f_1 - \hat{f}|}{f_1} + \frac{|f_2 - \hat{f}|}{f_2} \right] \quad (19)$$

10 where $\hat{f}$ is the estimate from clustering, single, or oracle closest in frequency to the true pitch values f_1 and f_2 .

Fig. 10a illustrates concurrent vowel pitch estimates 1010. Clustering values 1013, single values 1012, and true pitch values 1011 are shown. In Syn1 1010, the estimates rise at around 125-175 Hz and remain constant at around 150 Hz. In Sync2 1020, the estimates fall at around 250-200 Hz and rise at around 125-150 Hz. In Sync3 1030, these values rise at around 100-150 Hz and also at around 175-225 Hz. In sync 4 1040, the estimates fall at around 200-150 Hz and rise at around 150-200 Hz. Estimates start at around 50 (250) ms to remove edge effects.

Fig. 11a illustrates all-voiced mixture spectrogram with separate pitch trajectories for a male stating “Why were you away a year?” and a female stating “Nanny may know my meaning.” The first 250 ms and last 50 ms are excluded to remove edge effects in clustering due to initial and final silent regions. Figs. 11b-11d illustrate the values for truth 1110, oracle 1120 (Fig. 11b), clustering 1130 (Fig. 11c), and single 1140 (Fig. 11d).

Fig. 12 includes similar plots as Fig. 11 but with all-voiced mixture with crossing pitch trajectories. In Fig. 12c, the arrow 1250 denotes a “jump” due to differences in energy between sources in localized time-frequency regions.

Fig. 13 is a table that includes average $\%error$'s ($\%error_{avg}$) computed across time for examples studied in Figs. 10-12. For the synthetic concurrent-vowels task (Syn1-4, Fig. 10), GCT-based MRA provides accurate estimates under a variety of mixed pitch trajectories. The oracle estimates follow the true pitch values with

$\%error_{avg} < 0.04\%$ while the clustering scheme assigns pitch values across time for GCT-based MRA with $\%error_{avg} < 1.75\%$ (Fig. 13). The oracle and clustering of pitch candidates derived from GCT-based MRA exhibits lower $\%error_{avg}$ than single-region analysis in all cases.

5 For real speech, the oracle pitch values match truth with 0.00% average error in both separate and crossing conditions. Although close to truth for the separate case, it appears that median-based clustering is not optimal for exploiting the oracle candidates in the crossing case, with jumps in pitch values from distinct talkers (e.g., Fig. 12c, arrow 1250). This is likely due to the inability of the clustering method to
10 account for points in time in which one speaker is dominant in energy. Nonetheless, the accuracy of the oracle estimates demonstrates the feasibility of employing GCT-based MRA for multi-pitch estimation with an improved post-processing method. Finally, as in the synthetic cases, the oracle and clustering of the GCT-based MRA pitch candidates outperform the single-region method, thereby further illustrating the
15 benefits of exploiting multiple regions for analysis.

 Therefore, GCT-based MRA provides separability of pitch information for multi-pitch signals. Since the GCT can separate pitch information from multiple sources of similar energies, the assumption of a *single* dominant source does not need to be invoked when obtaining pitch candidates as is typically done for short-
20 time autocorrelation analysis methods (e.g., [12]). The accuracy of the pitch estimates obtained using GCT-based MRA on real and synthetic mixtures further demonstrates the feasibility of employing this analysis framework in conjunction with existing multi-pitch tracking techniques (e.g., using hidden Markov models [12]).

25

Multi-pitch Estimation Method 2: Pattern Recognition of GCT-based Features for Pitch Candidate Pruning

 Fig. 14 includes plot of candidate collections obtained across frequency regions from GCT analysis as described in the previous section "GCT-based
30 Analysis of Multi-pitch Signals". The candidate collections generally are located within a neighborhood of approximately 20 Hz for the pitch values and approximately 0.5 Hz/ms for the pitch derivatives. As shown in Fig. 14, a number

of time points exhibit candidates that are up to approximately 150 Hz away from the true pitch and approximately 1 Hz/ms from the true pitch derivative (e.g., at 250 ms). These substantial deviations highlight limitations in the underlying GCT analysis are used to obtain pitch candidates.

5 In relation to the multi-pitch estimation problem, certain example embodiments may apply an analysis scheme to obtain pitch. Specifically, one pitch candidate may be obtained from each time-frequency region of the short-time Fourier transform magnitude (STFTM) by first computing the GCT for the region and performing peak-picking. As an example, the STFTM on a mixture of two all-voiced sentences "Why were you away a year, Roy?" and "Nanny may know my
10 meaning" by two distinct speakers may be computed. To assess the value of these candidates, a collection of histogram slices computed for the pitch candidates obtained across all frequency regions for a single time point of analysis may be utilized. In addition, the two reference pitch tracks of the two speakers estimated
15 from their individual waveforms using a correlation-based pitch tracker may be employed [17].

Fig. 15 is a schematic that illustrates a method for obtaining pitch candidates according to certain example embodiments. In Fig. 15a, the STFTM of the mixture is computed and localized time-frequency regions are analyzed using the GCT (Fig.
20 15c) and a set of pitch candidates and features are obtained across frequency regions for each time point (Fig. 15c).

Fig. 16 illustrates a collection of histogram slices of pitch candidates obtained from GCT analysis (Fig. 16a) and a single histogram obtained at time = 64 ms. Spurious candidates obtained from GCT analysis can be seen across the
25 collection and in the individual slice (e.g., at 300 ~ 350 Hz, Fig. 16b).

Consistent with the single-speaker case, a number of spurious peaks can be obtained that are up to approximately 150 Hz away from the true pitch values of either speaker at time points across the mixture duration. For instance, at 64 ms, the histogram slice exhibits pitch candidates above 300 Hz while the true pitch values of
30 the two speakers are between 150 to 200 Hz (Fig. 16). In performing pitch tracking (e.g., using a state-space model) these errors are poorly modeled with simple distributions (e.g., Gaussian in the Kalman filtering framework), which may

motivated more sophisticated modeling (e.g., in a belief propagation framework). Nonetheless, as shown in Fig. 14, such an extension may also require the "noise" distributions surrounding the pitch candidates to vary with time (e.g., in Fig. 14, compare the variation of the pitch value at time = 250 ms with that at time = 75 ms).

5 To avoid this additional layer of complexity in modeling, certain example embodiments may consider an alternative approach to prune pitch candidates obtained from GCT analysis with the overall aim of improving multi-pitch tracking. Specifically, using characteristics of the GCT analysis itself, it may be determined whether a pitch candidate is spurious or not spurious. The value of a data-driven
10 approach may be assessed for peak selection.

Feature Extraction from the GCT

Fig. 17 demonstrates the GCT computed for a region in which the formant structure of one speaker in the mixture exhibits a formant peak. In Fig. 17a,
15 STFTM of mixture 1705 and localized region 1710 are analyzed. Fig. 17b demonstrates a zoomed-in version of the localized region 1710. In Fig. 17c, the GCT is computed without minimizing near-DC components. Fig. 17d illustrates the computed GCT with near-DC components minimized.

Fig. 18 demonstrates the GCT computed for a region in which the formant
20 structures of both speakers in the mixture exhibit formant peaks. In Fig. 18a, STFTM of mixture 1805 and localized region 1810 are analyzed. Fig. 18b demonstrates a zoomed-in version of the localized region 1810. In Fig. 18c, the GCT is computed without minimizing near-DC components. Fig. 18d illustrates the computed GCT with near-DC components minimized.

25 As shown in Fig. 17, the localized regions 1710 clearly contains harmonic structure from a speaker. The corresponding GCT exhibits a dominant set of peaks corresponding to these components. In contrast, in Fig. 18 demonstrates the GCT computed for a region in which the formant structure of both speakers has minimal amplitude. Specifically, Fig. 18d (e.g., GCT with near-DC components minimized
30 by 2-D gradient) illustrates that a large number of peaks are present with comparable amplitudes.

Various features may be extracted from the GCT in addition to the pitch and pitch-derivative candidates. The combined features can be used to determine whether the pitch candidate is spurious or non-spurious using pattern classification techniques. For example, the following features may be extracted from each
 5 localized time-frequency region (i.e., localized region) of the STFTM from GCT analysis:

1. pitch value obtained from location of dominant peak [18],
2. pitch derivative obtained from the mapping [18],
3. magnitude value of peak corresponding to pitch value and pitch value
 10 derivative,
4. normalized magnitude value of peak,
5. signal to noise ratio of a chosen peak,
6. DC value of localized region, and
7. overall energy of localized region.

15 A localized region centered at $n = n_0$ and $m = m_0$ may be defined as

$$s_w[n, m] == s[n - n_0, m - m_0]w[n, m] \quad (20)$$

where $s[n, m]$ denotes the full STFTM computed for the mixture signal and $w[n, m]$ denotes a 2-D Gaussian window. The GCT is defined as

$$S_w(\omega, \Omega) = FT[s_w[n, m]] \quad (21)$$

20 where FT denotes the 2-D Fourier transform.

Features 1-3 may be obtained from max peak-picking operation on the GCT magnitude:

$$Feature1 = f_0 = \frac{2\pi f_s}{N_{STFT} \omega_s \cos \theta} \quad (22)$$

$$Feature2 = \frac{\Delta f_0}{\Delta t} = \frac{f_0}{f_{center}} \tan \theta \quad (23)$$

25 $Feature3 = |S_w(\omega = \omega_s \sin \theta, \Omega = \omega_s \cos \theta)|. \quad (24)$

where N_{STFT} corresponds to the discrete-Fourier transform length used in computing the STFTM, f_s is the sampling frequency of the waveform, and f_{center} is the center frequency of the localized region. $\omega_s \cos \theta$ and $\omega_s \sin \theta$ correspond to the location of

the peak along the ω and Ω -axes of the GCT magnitude; θ is the angle between the peak location and the Ω -axis [18].

For Feature 4, the peak magnitude value is normalized by the sum of all magnitudes in the GCT:

$$5 \quad \text{Feature4} = \frac{\text{Feature3}}{\iint_{\omega, \Omega} |S_w(\omega, \Omega)|} \quad (25)$$

For Feature 5, a signal-to-noise ratio (SNR) is computed as

$$\text{Feature5} = 10 \log_{10} \frac{\sum_n \sum_m |K \cos(\Phi[n, m]) w[n, m]|^2}{\sum_n \sum_m |s[n - n_0, m - m_0] w[n, m]|^2 - \sum_n \sum_m |K \cos(\Phi[n, m]) w[n, m]|^2}.$$

(26) where $\Phi[n, m] = 2\pi\omega_s (n \cos \theta + m \sin \theta) + \varphi$ and $\Phi[n, m]$ is determined from the location of the GCT peak used in obtaining Features 1 and 2. The term K is similarly determined from the magnitude of this GCT peak [18, 21].

Finally, Features 6 and 7 are computed as

$$15 \quad \text{Feature6} = |S_w(\omega = 0, \Omega = 0)| \quad (27)$$

$$\text{Feature7} = \iint_{\omega, \Omega} |S_w(\omega, \Omega)|^2. \quad (28)$$

Pattern Recognition for Pitch Candidate Selection

Certain example embodiments provide for pitch candidate selection using pattern recognition techniques. A training set of data may be employed to obtain a transformation that allows for a simple hypothesis test to determine whether a pitch candidate is spurious or not spurious based on the features derived from GCT analysis. This transformation is then applied to the results of GCT analysis performed on a testing data set.

25 To obtain the transformation, features are generated using GCT analysis. Specifically, assuming that $\underline{x}_{t_i, f_k}$ is a feature vector obtained from a localized region in the k^{th} frequency band at time t_i (Figs. 14 and 16) and $\hat{f}_{0, i, k} = \underline{x}_{t_i, f_k}(1)$ denote its

corresponding pitch candidate, the collection of all $\underline{x}_{t_i, f_k}$ across the training set in time are grouped into spurious and not spurious candidates based on the absolute difference between the candidate pitch value and the true pitch value of either speaker. A feature vector $\underline{x}_{t_i, f_k}$ is labeled spurious (i.e., sp) if

$$\min\left(\left|\hat{f}_{0,i,k} - f_{01,i}\right|, \left|\hat{f}_{0,i,k} - f_{02,i}\right|\right) > \varepsilon \quad (29)$$

where $f_{01,i}$ and $f_{02,i}$ are the true pitch values of the two speakers at time i and ε is a threshold value. Otherwise, $\underline{x}_{t_i, f_k}$ is labeled not spurious (i.e., nsp). The term ε is chosen to be 20 Hz and is motivated from the observations regarding pitch candidates obtained from the GCT (e.g., Figs. 14 and 16).

A linear discriminative analysis (LDA) may be performed to obtain a transformation that maps $\underline{x}_{t_i, f_k}$ to a single value y_{t_i, f_k} to be used in a hypothesis testing/classification framework, i.e.,

$$y_{t_i, f_k} = \underline{w}^T \underline{x}_{t_i, f_k} \quad (30)$$

LDA obtains the transformation $\underline{w}$ that maximizes

$$J(\underline{w}) = \frac{|\tilde{m}_{sp} - \tilde{m}_{nsp}|^2}{\tilde{s}_{sp}^2 + \tilde{s}_{nsp}^2} \quad (31)$$

where $\tilde{m}_{sp}$, $\tilde{m}_{nsp}$, and $\tilde{s}_{sp}$, $\tilde{s}_{nsp}$ are the means and standard deviations of the transformed features (i.e., y_{t_i, f_k} 's) in the training data [22].

Fig. 19 includes results of LDA analysis of features on training data and maximum-likelihood fits for two classes (nsp 1910 and sp 1930). In Fig. 19, histograms computed for the two classes (nsp 1910 and sp 1930) of the transformed training data for a frequency band centered at 600 Hz are illustrated. In addition, the maximum-likelihood fits (ML_{NSP} 1920 and ML_{SP} 1940) to single Gaussians of these data are also illustrated. These fits 1920, 1940 are subsequently used in testing for classification using a hypothesis testing framework. Specifically, a feature vector $\underline{x}'_{t_i, f_k}$ obtained from the testing data is mapped to y'_{t_i, f_k} by $\underline{w}$, i.e.,

$$y'_{t_i, f_k} = \underline{w}^T \underline{x}'_{t_i, f_k} \quad (32)$$

y'_{t_i, f_k} is then classified as "not spurious" (nsp) if

$$\frac{p(y'_{t_i, f_k} | nsp_{f_k})}{p(y'_{t_i, f_k} | sp_{f_k})} > \frac{p(sp_{f_k})}{p(nsp_{f_k})} \quad (33)$$

and "spurious" (sp) if

$$\frac{p(y'_{t_i, f_k} | nsp_{f_k})}{p(y'_{t_i, f_k} | sp_{f_k})} < \frac{p(sp_{f_k})}{p(nsp_{f_k})} \quad (34)$$

to minimize the probability of error [22]. Here, $p(y'_{t_i, f_k} | nsp_{f_k})$ and $p(y'_{t_i, f_k} | sp_{f_k})$ are the frequency band-wise single-Gaussian fits obtained in training. Similarly, $p(sp_{f_k})$ and $p(nsp_{f_k})$ are the band-wise priors estimated from the training data.

Evaluation and Results

An example training set includes 40 mixtures of two all-voiced sentences ("Nanny may know my meaning" + "Why were you away a year, Roy?") spoken by distinct speakers. A similar set of 40 mixtures is obtained for the two all-voiced sentences in the testing set. However, speakers in the training set are distinct from those in the testing set. True pitch tracks for each individual speaker are obtained a priori using a correlation-based pitch estimator as in [18].

Fig. 20 includes plots obtained using a hypothesis testing framework on testing data. Specifically, Fig. 20 includes a plot of band-wise classification accuracy 2010 obtained using a hypothesis testing framework on testing data (i.e., the transformed features from GCT analysis) as well as plots of probability of falsely detecting a not-spurious candidate 2030 (false alarm), the probability of a miss 2040 (i.e., missing a "not-spurious" candidate), and the prior probability of a not-spurious peak 2020. As shown in Fig. 20, the candidate selection technique 2010 exhibits a classification accuracy exceeding 90%. Furthermore, inverting the probability of misses results in an approximate 80% detection probability.

Fig. 21a is an illustration of collected histogram slices of raw pitch candidates. Figs. 21(a) and 21(b) include raw pitch candidates and their corresponding histogram and Figs. 21(c) and 21(d) include pruned pitch candidates using collections of histogram slices. At $t = 60$, the true pitch values are about 156 Hz and 108 Hz. Observe that the pruned pitch candidates (Fig. 21c) exhibit fewer spurious candidates with substantially different (e.g., > 20 Hz absolute difference)

than the raw candidates (Fig. 21b). These results highlight the utility of the candidate selection method in improving the quality of candidates and is consistent with the empirical performance metrics of the hypothesis testing framework (Fig. 20). Furthermore, the results suggest that the pitch candidates selected from this method offer an improved basis for pitch estimation.

Clustering Techniques

To assess the value of the candidate selection method in multi-pitch estimation, certain example embodiments of the present invention may apply two simple clustering methods to generate pitch tracks of individual speakers. Candidates across both time and frequency regions are combined to form a set of candidates for a single time point and a simple median-based cluster is employed to select the pitch value at a point in time for both speakers. In the second method, a k-means clustering scheme is used. Both methods are applied to the pruned and raw pitch candidates.

To quantitatively assess the results, the pitch error metrics may be computed as in [8]. Specifically, a gross error may be defined as the condition where either assigned pitch value for a time point differs from the true pitch values by more than 20%:

$$\max \left(\frac{|\hat{f}_{0,i,k} - f_{01,i}|}{f_{01,i}}, \frac{|\hat{f}_{0,i,k} - f_{02,i}|}{f_{02,i}} \right) > 0.2 . \quad (35)$$

The total gross error (E_gross) is the percentage of time points across the entire mixture exhibiting a gross error. For time points in which there is no gross error, a fine error is computed based on the sum of percent errors between the two assigned pitch values and the true pitch values, i.e.,

$$FineError = 100 \left(\frac{|\hat{f}_{0,i,k} - f_{01,i}|}{f_{01,i}} + \frac{|\hat{f}_{0,i,k} - f_{02,i}|}{f_{02,i}} \right) \quad (36)$$

The total fine error is the average of fine errors across all time points in the mixture (E_fine) and the total error (E_total) is the sum of the total gross and fine errors, i.e.,

$$E_{total} = E_{fine} + E_{gross} \quad (37).$$

Fig. 22 illustrates an example of pruned and raw pitch candidates and resulting median-based clustering method to assign pitch candidates to time points.

Fig. 23 illustrates an example of pruned and raw pitch candidates and
5 resulting k-means clustering method to assign pitch candidates to time points.

For multi-pitch tracking, the median-based clustering that prunes candidates (Fig. 22a) provides higher fine and gross errors than median-based clustering using the raw candidates (Fig. 22b). In contrast, the k-means clustering results in smaller fine and gross errors for the pruned candidates (Fig. 23a) than clustering based on
10 the raw candidates (Fig. 23b).

Fig. 24 includes a table of average metrics across all test mixtures for the median clustering method. Specifically, average E_{gross} 2410, E_{fine} 2420, and E_{total} 2430 ($E_{total} = E_{fine} + E_{gross}$) between median-based clustering using pruned candidates 2440 versus clustering based on raw candidates 2450 are shown.
15 Similarly, Fig. 25 includes a table of average metrics across all test mixtures for the median clustering method. Specifically, average E_{gross} 2510, E_{fine} 2520, and E_{total} 2530 ($E_{total} = E_{fine} + E_{gross}$) between median-based clustering using pruned candidates 2540 versus clustering based on raw candidates 2550 are shown.

As shown in Figs. 24-25, the pruned candidates provide an approximate 9%
20 gross error reduction on average relative to the raw candidates for either clustering method. However, in k-means clustering, average fine errors are increased relative to median-based clustering (7.04% versus 6.45%). Overall, in this example, the best performing method appears to be k-means clustering of the pruned candidates with an average total error 12.62%, thereby demonstrating the potential value of the pitch
25 candidate pruning method.

Therefore, by extracting features motivated from observations of the GCT space and combining them with pattern classification methods, an improved set of pitch candidates can be obtained. Using simple clustering methods for pitch tracking, improvements in fine and gross errors in multi-pitch estimation can be
30 achieved using the results of pitch candidate selection.

Multi-Pitch Estimation using a Joint 2-D Representation of Pitch and Pitch Dynamics

The model of localized time-frequency region $s[n, m]$ (discrete-time and frequency n, m) of a narrowband STFT log-magnitude computed for a single voiced utterance. A simple model of the harmonic structure in $s[n, m]$ is a 2-D sinusoid resting on a DC pedestal:

$$s[n, m] \approx \alpha[n, m](1 + \sum_{k=1}^N \alpha_k \cos \phi_k [n, m]) \quad (38)$$

$$\Phi[n, m] \approx \omega_s(ncos\theta + msin\theta) + \varphi$$

where $\omega_s, \theta, \varphi$, and α correspond to the frequency, orientation, phase, and amplitude of the 2-D sinusoid, respectively. The GCT is the 2-D Fourier transform of $s[n, m]$ given by

$$S(\omega, \Omega) = 2\pi K \delta(\omega, \Omega) + 2\pi \delta(\omega + \omega_s \sin \theta, \Omega - \omega_s \cos \theta) + 2\pi \delta(\omega - \omega_s \sin \theta, \Omega - \omega_s \cos \theta). \quad (39)$$

Denoting f_{sp} as the waveform sampling frequency and N_{STFT} as the discrete-Fourier transform (DFT) length of STFT, the speaker's pitch f_0 at the center (in time) of $s[n, m]$ is related to $\omega_s \cos \theta$ through [10]

$$f_0 = \frac{2\pi f_{sp}}{N_{STFT}(\omega_s \cos \theta)} \quad (40)$$

A shift in f_0 (Δf_0) across a duration of Δn in $s[n, m]$ results in an absolute frequency shift of the k_{th} pitch harmonic by $k\Delta f_0$. Therefore, within $s[n, m]$:

$$\tan \theta \approx \frac{k\Delta f_0}{\Delta n} \quad (41).$$

Using f_0 from (40), k may be approximated as $k \approx f_{center}/f_0$ (shown in Fig. 1c(a)). The rate of change of f_0 ($\frac{\delta f_0}{\delta t}$) is:

$$\frac{\delta f_0}{\delta t} \triangleq \frac{\Delta f_0}{\Delta n} = \frac{f_0 \tan \theta}{k f_0} \quad (42).$$

Extending the model of (38) to the condition of N speakers in $s[n, m]$, the localized time-frequency region may be approximated as:

$$s[n, m] \approx \sum_{i=1}^N [K_i + \alpha_i \cos(\Phi_i[n, m])] \quad (43)$$

$$\Phi[n, m] = \omega_s(ncos\theta + msin\theta) + \varphi .$$

Such that the GCT is:

$$S(\omega, \Omega) = 2\pi \sum_{i=1}^N K_i \delta(\omega, \Omega) + \pi \sum_{i=1}^N K_i \delta(\omega + \omega_s \sin \theta, \Omega - \omega_s \cos \theta) \\ + \pi \sum_{i=1}^N K_i \delta(\omega - \omega_s \sin \theta, \Omega + \omega_s \cos \theta) \quad (44)$$

The GCT therefore *jointly* represents both pitch and pitch derivative information distinctly for each speaker.

Fig. 26A is an illustration of a multi-pitch estimation method 2600 for a mixed waveform 2610 that employs clustering and Kalman filtering. The multi-pitch estimation method 2600 includes short-time analysis 2620, GCT analysis 2630, discriminate-based pitch candidate pruning 2640, and a clustering and Kalman filtering framework 2650.

Mixture waveforms 2610 may be analyzed using the short-Time Fourier transform (STFT) 2620 to form the log spectrogram. In certain embodiments, a 32-ms Hamming window, 1-ms frame interval, and 512-point discrete Fourier transform (DFT) may be used to compute the STFT, denoted as log-STFTM. A representative log-STFTM 2670 computed for a mixture of the “Walla Walla” and “Lawyer” sentences spoken by two female speakers is shown in Fig. 26B.

Specifically, the “lawyer” and “Walla Walla” sentences were selected from sentences s5 and s6 of the following sentences:

s1 – “May we all learn a yellow lion roar.”

s2 – “Why were you away a year, Roy?”

s3 – “Nanny may know my meaning.”

s4 – “I’ll willingly marry Marilyn.”

s5 – “Our lawyer will allow your rule.”

s6 – “We were away in Walla Walla.”

s7 – “When we mow our lawn all year.”

s8 – “Were you weary all along?”

STFTM results may subsequently be used for GCT analysis 2630. A 2-D high-pass filter was applied to log-STFTM to reduce the effects of the DC components in the GCT representation and is denoted as STFTM_{HP} [6]. Localized

regions of size 800 Hz by 100 ms may be extracted using a 2-D Hamming window from the magnitude of both log-STFTM and log-STFTM_{HP}. Overlap factors of 10 and 4 may be used along the time and frequency dimensions and result in a set of center frequencies for GCT analysis along the frequency axis and overlapped
 5 regions for analysis in time. A 2-D DFT of size 512 by 512 may be used to compute two GCT's: GCT_M (from log-STFTM) and GCT_{HP} (STFTM_{HP}). Seven features may be extracted:

1. Pitch estimate $\hat{f}_0$ from peak-picking the dominant peak in the magnitude of GCT_{HP} ($|GCT_{HP}|$) and Equation (40).
- 10 2. Pitch-derivative estimate $\partial \hat{f}_0 / \partial t$ obtained from Equation (42) from the dominant peak in $|GCT_{HP}|$
3. Amplitude of the dominant peak in $|GCT_{HP}|$
4. Normalized value of the amplitude of the dominant peak in $|GCT_{HP}|$
5. "Harmonic to noise ratio" of dominant peak in $|GCT_{HP}|$
- 15 6. DC value of GCT_M
7. Overall energy of GCT_M

Feature 4 may be computed as:

$$Feature\ 4 = \frac{Feature\ 3}{\iint |GCT_{HP}(\omega, \Omega)|} \quad (45)$$

Feature 5 may be computed as:

$$Feature\ 5 = 10 \log_{10} \frac{\sum_n \sum_m |\alpha \cos(\Phi[n, m])|^2}{\sum_n \sum_m |s[n, m]|^2 - \sum_n \sum_m |\alpha \cos(\Phi[n, m])|^2}, \quad (46)$$

20 where $\alpha = Feature\ 3$ and $s[n, m]$ corresponds to the localized region of log-STFTM.

Features 3 - 7 relate to properties of the GCT not captured by the pitch and pitch-derivative and are used in pitch candidate pruning 2640.

GCT analysis of multi-pitch signals may result in $\hat{f}_0$ far removed from the
 25 true pitch value (denoted as f_0), presumably from regions in which harmonic structure exhibits low amplitudes or substantial overlap from multiple speakers [19]. To account for these "spurious" candidates, linear discriminate analysis (LDA) [22] may be applied to the previously described set of 7 features to prune the candidates (LD-based pruning 2640). In training, certain embodiments define a "spurious"

candidate as one in which $|\hat{f}_0 - f_0| > \delta$. The term δ is set to, 3σ where σ is the standard-deviation of the one-step differences in the pitch values of the training data. In one example embodiment $\sigma = 4.85$ Hz. A discriminate function is trained for each center frequency in GCT analysis and applied in a band-wise fashion to prune
 5 the candidates.

Fig. 26C is an illustration of band-wise classification performance of linear discriminate analysis on test data per band. Fig. 26D is an illustration of resulting binary mask of pruning of the plot shown in Fig. 27A with 1's and 0's. The pruning is performed across time and center frequency for the mixture in Fig. 26B. In Fig.
 10 26D, 1's denote regions in which the candidate is kept while 0's denote regions in which they are discarded.

Given the pruned candidates across time-frequency regions, k-means clustering 2650 may be used to obtain local estimates in time. As our mixtures contained all-voiced speech from two speakers, two centroids were extracted from
 15 pruned candidates across all frequency bands at each time point. Specifically, certain embodiments may perform clustering along *both* the pitch and pitch-derivative dimensions, where the pitch derivative estimate is that tied to the pruned pitch value (i.e., Feature 2). Therefore, such embodiments account for conditions where pitch values may be identical but pitch-derivatives may differ for two
 20 speakers. To generate the pitch track for each speaker, each pair of centroids at a point in time may be used as observations to a pair of Kalman filters (KF) [24]. For each speaker i , certain embodiments may adopt a state-space model

$$x_{t+1,i} = Ax(t+1,i) + v_t \quad (47)$$

$$y_{t,i} = \underline{x}_{t,i} + \underline{w}_t$$

Where $A = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix}$, $\underline{y}_{t,i} = \begin{bmatrix} \tilde{f}_0(t,i) \\ \partial \tilde{f}_0 / \partial t(t,i) \end{bmatrix}$ is the centroids, and $\underline{x}_{t,i} =$
 25 $\begin{bmatrix} f_0(t,i) \\ \partial f_0 / \partial t(t,i) \end{bmatrix}$ is the true state. The terms $\underline{v}_t$ and $\underline{w}_t$ are Gaussian noise terms [24].

Given the assignment of a centroid to a speaker, the standard KF equations are used to generate the pitch track. In training, the covariances of $\underline{v}_t$ and $\underline{w}_t$ are obtained using the optimal assignment of centroids on a training set of mixtures. Optimal is

defined as when the observation is closest in normalized distance to the true state, where normalization is done by the means and standard deviations of the pruned candidates at each time point.

To perform assignment of the centroids to each speaker/pitch track in testing, certain embodiments compute distances between the predicted states of the two pitch tracks $\hat{x}_{t,1|t-1}$ and $\hat{x}_{t,2|t-1}$ and the two observations $\underline{y}_{t,a}$ and $\underline{y}_{t,b}$ at time t .

Specifically, $\chi_{1,a}$ may be defined as:

$$\chi_{1,a} = (\underline{y}_{t,a} - \hat{x}_{t,1|t-1})^T \Lambda_{t|t-1}^{-1} (\underline{y}_{t,a} - \hat{x}_{t,1|t-1}) \quad (48)$$

10

Where $\Lambda_{t|t-1}^{-1}$ is the covariance associated with the prediction at time t .

$\chi_{1,b}$, $\chi_{2,a}$, and $\chi_{1,b}$ may be similarly defined. For $\hat{x}_{t,1|t-1}$, the minimum of $\chi_{1,a}$ and $\chi_{2,a}$ may be used to make the assignment to the corresponding observation, the same rule may be applied for $\hat{x}_{t,2|t-1}$ but with $\chi_{1,b}$ and $\chi_{2,a}$. If

15 $\hat{x}_{t,1|t-1}$ and $\hat{x}_{t,2|t-1}$ acquire the same observation (e.g. if they both acquire $\underline{y}_{t,a}$), the assignments are changed based on the following criterion:

$$\text{If } (\chi_{1,a} + \chi_{2,a} > \chi_{2,a} + \chi_{1,a}) \text{ assign } \underline{y}_{t,b} \text{ to } \hat{x}_{t,2|t-1} \quad (49)$$

Otherwise, assign $\underline{y}_{t,b}$ to $\hat{x}_{t,1|t-1}$

20 The same rule is applied if $\hat{x}_{t,1|t-1}$ and $\hat{x}_{t,2|t-1}$ both acquire $\underline{y}_{t,b}$ but with $\underline{y}_{t,a}$ replacing $\underline{y}_{t,b}$ in (49). This assignment uses individual uncertainties of predicted observations and the combined uncertainty of both assignments to prevent pitch tracks from merging. Fixed-interval smoothing (across the entire duration of the pitch track) is applied to the filtered estimates [24]. Utilizing both pitch and
25 pitch-derivative information in multi-pitch estimation hereinafter is being referred as “ $f_0 - \partial f_0 / \partial t$.”

To assess the utility of the GCT's joint representation of pitch and pitch-dynamics, certain embodiments use a reference system that does not utilize $\partial f_0 / \partial t$ in estimation. The candidates are pruned based on the $|\hat{f}_0 - f_0| > \delta$ criterion, but k-means clustering is done using only the pitch values. In tracking, the state-space
30

model of (47) is modified such that $A = 1$, $y_{t,i} = \tilde{f}_0(t, i)$ is the centroid, and $x_{t,i} = f_0(t, i)$. This approach is referenced hereinafter as “ f_0 _only.”

For evaluation, an example data set, collected consisting of 8 males (m1 - m8) and 8 females (f1 - f8) speaking 8 all-voiced utterances (sentences s1-s8, outlined above) and sampled at 16 kHz was used. Data was obtained from speakers that maintained voicing throughout each utterance. Reference (or "true") pitch values of the sentences were obtained using Wavesurfer prior to mixing [25]. Speech files were pre-emphasized at a 0-dB overall signal-to-signal ratio. To train the LDA-based pruning and Kalman filters 2650, mixtures generated from first four male (m1-m4) and first four female (f1-f4) speakers, speaking sentences s1 - s4 were used. In testing, mixtures generated from second four male (m5- m8) and second female (f5 - f8) speakers speaking sentences s5 - s8 were used. Distinct speakers and sentences were used in each mixture such that train and test sets consisted of 336 total mixtures each.

The test data was then divided into mixtures of “separate” and “close” pitch track conditions, with “close” referring to mixtures where at least one time point contains a pair of pitch values within 10 Hz of each other. This accounts for 136 mixtures, the majority of which contained either crossings, or both crossings and mergings. The remaining 200 mixtures are considered separate. Representative mixtures are shown in Figs. 26F-26K.

In order to obtain a quantitative metric for performance, certain embodiments of the present invention define a root-mean-squared-errors (RMSE) as:

$$\text{RMSE} = \sqrt{\frac{1}{2L} \sum_{i=1}^2 \sum_{t=1}^L (\hat{x}_{t,i} - x_{t,i})^2} \quad (50)$$

where L is the length of the mixture and $\hat{x}_{t,i}$ and $x_{t,i}$ are the reference and estimated pitch values, respectively.

Fig. 26E shows average RMSEs in both the “separate” and “crossing” datasets for the two described estimation methods along with standard errors. Figs. 26F-26G show the results of a “separate” case comparing the $f_0 - \partial f_0 / \partial t$ and f_0 _only approaches. Consistent with the quantitative results in Fig. 26E, Figs. 26F-26G exhibit similar performance to f_0 _only and obtains reasonable

estimates of pitch values. In contrast, the $f_0 - \partial f_0 / \partial t$ approach outperforms f_{0_only} by approximately 6 RMSE in the *close* conditions (Figs. 26H-26K). This is due to a more accurate estimate of the trajectories when they exhibit crossings (e.g., compare in Figs. 26H-26I, approximately 1000 ms). Nevertheless, an outstanding
 5 limitation of $f_0 - \partial f_0 / \partial t$ is when pitch tracks exhibit similar pitch values and pitch-derivatives. As an example, observe that erroneous estimates are made by $f_0 - \partial f_0 / \partial t$ near 150 ms in Figs. 26J-26K, where the two pitch tracks are close in absolute frequency and have similar slopes.

10 Analysis and Synthesis of Speech Using 2-D Sinusoidal Series

Example embodiments of the present invention may model a localized time-frequency region $s[n, m]$ of a narrowband short-time Fourier Transform magnitude based on a sinusoidal series modulated by an envelope:

$$15 \quad (51) \quad s[n, m] \approx \alpha[n, m](1 + \sum_{k=1}^N \alpha_k \cos \phi_k[n, m])$$

$$(52) \quad \phi_k[n, m] = \omega_s k(ncos\theta + msin\theta) + \psi_k$$

Accordingly, a series of N harmonically related sinusoids with spatial frequencies $k\omega_s$, orientation θ , amplitudes α_k , and phases ψ_k resting on a DC
 20 pedestal modulates a slowly-varying envelope $\alpha[n, m]$. The GCT is the 2-D Fourier transform of $s[n, m]$:

$$S(\omega, \Omega) = A(\omega, \Omega) + .05 \sum_{k=1}^N e^{-j\psi_k} A(\omega + \kappa\omega_s \sin \theta, \Omega - \kappa\omega_s \cos \theta) + e^{j\psi_k} A(\omega - \kappa\omega_s \sin \theta, \Omega + \kappa\omega_s \cos \theta)$$

where ω and Ω map to n and m , respectively (Fig. 26N). An approximate version of this model using a single sinusoidal carrier was derived in [21, 25] from properties of periodic source signals in speech. The sinusoidal series model may
 25 also be used to perform speaker separation as in the single sinusoid model. The representation presented in sinusoidal series model may further be exploited to perform multi-pitch analysis.

Certain embodiments of the present invention remove the approximation to account for multiple harmonically related carriers as observed in the GCT [25]. For voiced speech, formant structures may be mapped to the near-DC region of the GCT (Fig. 26N). The parameters of $\phi_k[n, m]$ for $k = 1$ relate to pitch and pitch-dynamic information while $k = 2, 3, \dots, H$ are harmonically related to this carrier.

Certain embodiments extend this model to account for onsets/offsets (e.g., vertical edges, Fig. 26N) in the spectrogram as they map *along* the GCT ω -axis from image processing principles and as observed in [20].

To account for noisy/unvoiced regions, certain embodiments consider the short-time Fourier transform of a white Gaussian process and invoke an independence assumption between each time-frequency unit of the spectrogram. These embodiments model the magnitude of each unit as an independent realization of a 2-D Rayleigh process $P_r[n, m]$ as in [26]. Denoting σ_R as the parameter of the Rayleigh distribution, the 2-D autocorrelation function and GCT-based power spectral densities are then

$$\begin{aligned} R_{pp}[\iota, \kappa] &= E[P_r[n, m]P_r[n + \kappa, m + \iota]] \\ &= 2\sigma_R^2, \kappa = 0, \iota = 0 \\ &= \frac{\pi}{2}\sigma_R^2, \text{otherwise} \end{aligned} \quad (54)$$

$$Spp(\omega, \Omega) = \sigma_R^2 \left(\frac{\pi}{2} \delta(\omega, \Omega) + \frac{4-\pi}{2} \right) \quad (55)$$

From (55), noise content within local time-frequency regions, on average, is spread to all regions of the GCT. To account for this behavior in individual regions, certain embodiments adopt the model of (51) and extract distinct carrier positions for each region.

The log of $s[n, m]$ may also be considered:

$$S_{log}[n, m] = \log \alpha[n, m] + \log(1 + \sum_{k=1}^n \beta_k \cos \phi_k[n, m]) \quad (56)$$

$$S_{log}[n, m] \approx K + \sum_{k=1}^n \beta_k \cos \phi_k[n, m] \quad (57)$$

From (57), example embodiments approximate $\log \alpha[n, m]$ as a DC constant k based on observations that the log tends to “flatten” the underlying 2-D envelope in localized regions, thereby allowing for improved estimation of the 2-D carrier

frequencies [19]. Moreover, $\log(1 + \sum_{k=1}^n \beta_k \cos \phi_k [n, m])$ is periodic with a fundamental spatial frequency ω_s as in (51) since the log operation maintains the periodicity of its argument such that $\phi_k [n, m]$ are defined as in (52), but β_k are arbitrary amplitudes distinct from α_k .

5 Figs. 26L-26O illustrate that certain embodiments of the present invention may employ GCT to represent information across the entire space in a distributed manner. Specifically, Fig. 26L illustrates a narrowband spectrogram 26010 having a local region 26020. Fig. 26 M illustrates a zoomed-in portion of the local region 26020 shown in Fig. 26L. Fig. 26N illustrates the GCT representation with near-DC
10 component to be removed for analysis and synthesis 26030 and series modulated components 26040. Fig. 26O illustrates demodulation 26050 in the GCT domain for analysis and synthesis using series to reconstruct near-DC terms 26060.

Fixed Region Size Analysis/Synthesis

15 For analysis/synthesis using the described model, example embodiments adopt the experimental setup of [21] in which example embodiments first remove the $A(\omega, \Omega)$ component in the GCT. This is based on the observations that interference in the GCT domain (e.g., other speakers) tends to be concentrated at the origin. Certain embodiments may recover $A(\omega, \Omega)$ using series-based sinusoidal
20 demodulation (Fig. 26N).

A narrowband magnitude spectrogram $s_{full}[n, m]$ and log spectrogram $s_{log-full}[n, m] = \log s_{full}[n, m]$ may be computed for the signal. A high-pass 2-D filter may further applied to both spectrograms to remove near-DC terms for GCT analysis. The filtered results are denoted herein as $s_{HP}[n, m]$ and $s_{log-HP}[n, m]$.

25 Localized regions are extracted from $s_{full}[n, m]$, $s_{HP}[n, m]$, and $s_{log-HP}[n, m]$ using a 2-D Hamming window denoted as $s[n, m]$, $s_{local, hp}[n, m]$, $s_{local, log-hp}[n, m]$. A GCT is computed for $s_{local, log-hp}[n, m]$ using a 2-D discrete Fourier Transform (GCT). Peak-picking in the GCT domain is then used to estimate the $k = 1$ carrier parameters ω_s , θ , and ψ_1 .

30 A 2-D carrier $\cos \hat{\phi}_1[n, m]$ may be generated from these parameters and multiplied by $s[n, m]$. This result may be low-pass filtered to obtain a scaled

estimate of $a[n, m]$, denoted as $\hat{a}[n, m]$. Subsequently, additional carriers are generated by scaling ω_s by $k = 2, 3, \dots, N$, where N is such that $N\omega_s < \pi$ and extracting ψ_k values at the carrier locations in the GCT. As in the $k = 1$ case, each carrier is multiplied by $s[n, m]$ and low-pass filtered to obtain an estimate of
 5 $a[n, m](\hat{a}[n, m])$.

Least-squares error (LSE) fitting is used to solve for the set of gain parameters γ_k by setting the known model components of (51) to $s[n, m]$:

$$S[n, m] = K + \sum_{k=1}^N \gamma_k \hat{a}_k[n, m] + \sum_{k=1}^N \hat{a}_k[n, m] \cos \phi_k[n, m] \quad (58)$$

10

The reconstructed spectrogram may be computed using 2-D overlap-add. The sinusoidal series model of embodiments of the present invention may also be used to solve for multiple speakers (as in [21]).

15 Adaptive Analysis and Synthesis

Certain embodiments of the present invention *adaptively* modify region sizes for the GCT, inspired by evidence in mammalian auditory studies of adaptive signal processing mechanisms that adapt to properties of the analyzed signal itself [20]. Specifically, these embodiments adapt region sizes based on a quantitative metric
 20 that assesses the relative “salience” of the proposed signal model in each localized region, thereby allowing for distinct resolutions of the GCT analysis based on the signal analyzed.

Figs. 27A-27D includes four spectrograms of synthetic signals. Synth1 (Fig. 27A) is a pulse_train with rising pitch (150 to 200 Hz). Synth2 (Fig. 27B) is a pulse
 25 train with_fixed pitch of 200 Hz. Synth3 (Fig. 27C) is a Gaussian white noise, and Synth4 (Fig. 27D) is a single impulse. Synth1 and Synth2 both excite a formant structure. In one embodiment, the formant structure may contain formant frequencies (bandwidths) of 669, 2349, 2972, 3500 Hz (65, 90, 156, 200 Hz).

Certain embodiments of the present invention may perform
 30 analysis/synthesis on these waveforms using the 2-D modeling developed in Equations (51)-(58), across distinct sets of fixed regions sizes. These fixed sizes

may be varied in time ranging from 20 to 50 ms in 2-ms steps while in frequency from 625 Hz to 1000 Hz in 62.5-Hz steps. As a quantitative metric for comparison, the global signal-to-noise ratio (SNR) between the original and re-synthesized waveforms may be used.

5 Fig. 28 is a table that outlines the SNR value for all four signals shown in Fig. 27A-27D across four distinct regions sizes. These sizes correspond to those that maximized the SNR for each distinct signal such that the diagonal of the table are the maximal SNR values in re-synthesis. From these results, it can be observed that four distinct sizes are obtained for each signal. Furthermore, a sub-optimal
10 selection of the region size for distinct signals can result in substantial SNR degradations (e.g., the optimal size for Synth2 leads to an SNR of 3.86 dB for Synth4, ~15 dB below the optimal size for Synth4). These results demonstrate that a “fixed tiling” of 2-D space exhibits limitations in reconstruction based on distinct properties of the signal itself.

15 Certain embodiments of the present invention may employ a “relative salience” of 2-D carrier frequencies in the GCT with respect to the rest of the GCT content. This metric quantitatively assesses the extent to which the series-based 2-D amplitude model is valid for a given region and may be used to guide adaptive region-growing and selection.

20 Let $s_{log}[n, m](s_{log-hp}[n, m])$ denote a local region of the (high-pass filtered) narrowband log-spectrogram for a given signal such that its corresponding GCT is $S_{log}(\omega, \Omega)(S_{log-hp}(\omega_d, \Omega_d))$. As discussed in Section 2, extracting the dominant peak magnitude of the $S_{log-hp}(\omega, \Omega)(|S_{log-hp}(\omega_d, \Omega_d)|)$ can be used to derive the carrier parameters ω_s , θ , and ψ_1 of a sinusoid denoted as $c_1[n, m]$. The
25 term $c_1[n, m]$ is scaled such that its GCT magnitude has a dominant peak value of $|S_{log-hp}(\omega_d, \Omega_d)|$. The remaining harmonically related carriers $c_k[n, m]$, $k = 2, 3, \dots, N$ are then obtained by scaling the parameters of the dominant carrier.

A “salience ratio” (SR) as the ratio of following energies may be obtained as:

$$SR = 10 \log_{10} \frac{E_c}{\iint_{\omega, \Omega} |S_{log}(\omega, \Omega)|^2 d\omega d\Omega - E_c} \quad (59)$$

$$30 \quad E_c = \iint_{\omega, \Omega} \sum_{k=1}^N |c_k(\omega, \Omega)|^2 d\omega d\Omega \quad (60)$$

In the above terms, the denominator is the energy difference in the local region of the original (non-filtered) narrowband spectrogram and the carriers. This metric relates the relative energy contributions of the carrier positions in the signal model to the overall region analyzed.

5 To adapt and select region sizes based on SR, certain embodiments first perform GCT analysis across the spectrogram of the signal analyzed using a fixed region size with a modified 2-D Hamming window that satisfies the constant overlap-add property [11]. This is referred to as base tiling. In each region, the SR metric, shown in Equation (59) may be computed. The result of this initial analysis
10 is a 2-D grid of SR values (Fig. 29A). Each base region is grown by examining its neighbors' SR values. Specifically, in Figs. 29A-29B SR_{base} denotes the SR value of a base region with neighboring SR_{top} , SR_{bottom} , SR_{left} , SR_{right} as shown. Furthermore, $SR_{merged,neighbor}$ denotes the SR value computed using the combined windows of the base and one of its neighboring regions. The base region
15 is recursively merged with its neighbors based on:

- a. Compute $SR_{merged,neighbor}$ of the base region for all of its neighbors (top, bottom, left, right)
- b. Determine the maximum of the four SR values computed in **a.** (denoted as $SR_{merged,max}$) with its corresponding neighbor $max_neighbor$.
- 20 c. If $SR_{merged,max} < \max(SR_{base}, SR_{top}, SR_{bottom}, SR_{left}, SR_{right})$ terminate the algorithm by creating a new region SR_{merged} equal to the base region. Otherwise, merge base region with $max_neighbor$ to form SR_{merged} with corresponding SR value $SR_{merged,max}$. Determine the new neighbors of SR_{merged} by its four edges. Use SR_{merged} as the *base* region in A1) to complete
25 the recursion.

The algorithm iteratively grows each base region until the SR value of any resulting merged region is less than that of the unmerged region. The order of the base regions merged is based ordering the SR values of all base regions in descending order. In case the neighbor of a base region has already been
30 incorporated into a previously merged region, it is excluded from the SR computations and comparison in the algorithm. The neighbors of any region are strictly those along its four vertical edges such that only rectangular regions are

grown (Fig. 29B). After all base regions have been processed, the resulting set of merged regions is used in 2-D demodulation for

Analysis and synthesis. The 2-D Hamming windows in each merged region are summed and used to extract the appropriate coordinates of the spectrogram and demodulated individually. The results are summed across all merged regions to reconstruct the spectrogram. The constant-overlap-add property of the 2-D Hamming windows guarantees a unity system if demodulation is not performed. The reconstructed spectrogram is combined with the phase of the original spectrogram and inverted

for waveform reconstruction.

Full Audio Source Separation System

Figs. 30A and 30B are a high-level illustration of an audio separation system according to certain example embodiments. Two audio source output audio signals $s_1[n]$ 3010 and $s_2[n]$ 3020. These speaker signals 3010, 3020 include distinct pitch trajectories $f_1[n]$ 3015 and $f_2[n]$ 3025. The speaker signals $s_1[n]$ 3010 and $s_2[n]$ 3020 are presented to a two-dimensional multi-pitch estimation subsystem 3040 as a mixed signal $y[n]$ 3030. The two-dimensional multi-pitch estimation subsystem 3040 determines pitch estimates $\hat{f}_1[n]$ 3055 and $\hat{f}_2[n]$ 3060 of pitch values $f_1[n]$ 3015 and $f_2[n]$ 3025. The pitch estimates 3055, 3060 are applied to a source separation subsystem 3070. The source separation subsystem 3070 determines estimates $\hat{s}_1[n]$ 3080 and $\hat{s}_2[n]$ 3090 of source signals $s_1[n]$ 3010 and $s_2[n]$ 3020.

Figs. 31A and 31B are a high-level illustration of a two-dimensional multi-pitch estimation subsystem 3040 according to certain example embodiments. The mixture signal $y[n]$ 3030 is presented to the system for spectrogram computation 3110 to compute the STFTM of the mixture 3150. Localized time-frequency regions are analyzed using the GCT 3170 (localized GCT processing 3120). In pitch and feature extraction 3130, one pitch candidate may be obtained from each time-frequency region 3160 of the short-time Fourier transform magnitude (STFTM) by first computing the GCT 3170 for the region and performing peak-picking to obtain

a set of pitch candidates across frequency regions for each time point 3180. Pitch value assignment 3140 is performed to determine whether a pitch candidate is spurious or not spurious. Based on the results obtained from pitch value assignment 3140, pitch estimates $\hat{f}_1[n]$ 3055 and $\hat{f}_2[n]$ 3060 of pitch values $f_1[n]$ 3015 and $f_2[n]$ 3025 are obtained.

Figs. 32A and 32B are a high-level illustration of a two-dimensional source separation subsystem according to certain example embodiments. The pitch estimates $\hat{f}_1[n]$ 3055 and $\hat{f}_2[n]$ 3060 are used to obtain pitch and pitch-derivative candidates 3210, 3215. A 2-D AM model 3220, 3225 is applied to determine the speech contents in the resulting signals 3211, 3216. The signals 3221, 3226 are summed and a least-squared error fit is employed to estimate a gains parameter 3240. A 2-D AM model 3250, 3255 is applied to determine the speech contents and spectrogram reconstruction is performed 3260 by overlap-add (OLA) of localized time-frequency regions 3150. The 2-D source separation system 3070 determines estimates $\hat{s}_1[n]$ 3080 and $\hat{s}_2[n]$ 3090 of speaker signals $s_1[n]$ 3010 and $s_2[n]$ 3020 using the reconstructed spectrogram 3280.

Fig. 33 is a high-level illustration of a system for processing an acoustic signal 3300 according to example embodiments of the present invention. A frequency-related representation preparer 3310 prepares a first frequency-related representation of the acoustic signal over time 3320. A two-dimensional transformer 3330 computes a two-dimensional transform of plural two-dimensional localized regions 3340 of the first frequency-related representation, each less than an entire frequency range of the first frequency related representation, to provide a two-dimensional compressed frequency-related representation with respect to each two-dimensional localized region. An identifier 3350 identifies at least one pitch for each of the plural regions 3360. A processor 3370 processes the pitch from the plural regions to provide multiple pitch estimates over time 3380.

Fig. 34 is a high-level illustration of a system for processing a mixed signal 3400 according to example embodiments of the present invention. A localizer 3410 localizes multiple time-frequency regions of a spectrogram of the mixed signal to obtain one or more acoustic properties 3420. A pitch estimate provider 3430

provides a separate pitch estimate of each of the multiple signals 3440 at a time point as a function of combining the one or more acoustic properties. A signal recoverer 3450 recovers at least one of the multiple signals 3460 as a function of the separate pitch estimate.

5 It should be understood that procedures, such as those illustrated by flow diagrams or block diagrams herein or otherwise described herein, may be implemented in the form of hardware, firmware, or software. If implemented in software, the software may be implemented in any software language consistent with the teachings herein and may be stored on any computer readable medium
10 known or later developed in the art. The software, typically, in form of instructions, can be coded and executed by a processor in a manner understood in the art.

 While this invention has been particularly shown and described with references to example embodiments thereof, it will be understood by those skilled in the art that various changes in form and details may be made therein without
15 departing from the scope of the invention encompassed by the appended claims.

References, all of which are incorporated by reference in their entirety:

- [1] D. Morgan, E. George, L. Lee, and S. Kay, "Cochannel Speaker Separation by Harmonic Enhancement and Suppression," IEEE TSAP, v5, pp. 407-424, 1997.
- 5 [2] T. Quatieri and R. Danisewicz, "An Approach to Co-channel Talker Interference Suppression Using a Sinusoidal Model for Speech," IEEE TASSP, v38, 1990.
- [3] S. Schimmel, L. Atlas, K. Nie, "Feasibility of Single Channel Speaker Separation Based on Modulation Frequency Analysis," ICASSP 2007, Honolulu, HI, USA.
- [4] G. Hu and D. Wang, "Monaural speech segregation based on pitch tracking and
10 amplitude modulation," IEEE TNN, vol. 15, 2004.
- [5] T. Quatieri, "2-D Processing of Speech with Application to Pitch Estimation," ICSLP 2002.
- [6] T. Ezzat, J. Bouvrie, and T. Poggio, "Spectrotemporal Analysis of Speech Using 2-D Gabor Filters," Interspeech 2007, Antwerp, Belgium.
- 15 [7] T. Ezzat, J. Bouvrie, and T. Poggio, "AM-FM Demodulation of Spectrograms Using Localized 2-D Max-Gabor Analysis," ICASSP, 2007.
- [8] T. Wang and T. Quatieri, "Exploiting Temporal Change of Pitch in Formant Estimation," ICASSP, 2008, Las Vegas, NV, USA.
- [9] T. Chi, P. Ru, S. Shamma, "Multiresolution Spectrotemporal Analysis of
20 Complex Sounds," JASA v118, pp. 887 - 906, 2005
- [10] T. Wang and T. Quatieri, "2-D Processing of Speech for Multi-Pitch Analysis," Interspeech 2009, Brighton, UK.
- [11] T. Quatieri, Discrete-time Speech Signal Processing: Principles and Practice. Upper Saddle River, NJ: Prentice-Hall, 2001.

- [12] M. Wu, D. Wang, and G. Brown, "A Multipitch Tracking Algorithm for Noisy Speech," *IEEE TASL*, v11, pp. 229 - 241, 2003.
- [13] U.S. Patent No. 7,574,352 to Thomas F. Quatieri, Jr.
- [14] T. Tolonen and M. Karjalainen., "A Computationally Efficient Multi-pitch
5 Analysis Model, " *IEEE Trans. on Speech and Audio Proc.*, 8:708-716, 2000.
- [15] T. Ezzat, J. Bouvrie, and T. Poggio, "Spectro-temporal Analysis of Speech Using 2-D Gabor Filters," *ISCA Interspeech*, 2007.
- [16] K. N. Stevens, *Acoustic Phonetics*. Cambridge, MA: MIT Press, 1998.
- [17] Y. Medan, E. Yair, and D. Chazan, "Super Resolution Pitch Determination of
10 Speech Signals," *IEEE Trans. Sig. Proc.*, 39:40-48, 1991.
- [18] "Program 1306 Quarterly Report," MIT Lincoln Laboratory, January 1st, 2009 - March 31, 2009.
- [19] T. T. Wang and T. F. Quatieri, "2-D Processing of Speech for Multi-Pitch Analysis," *Interspeech 2009*, Brighton, UK, 2009, *to be presented*.
- 15 [20] T. Ezzat, J. Bouvrie, and T. Poggio, "Spectrotemporal Analysis of Speech Using 2-D Gabor Filters," presented at International Conference on Spoken Language Processing, Antwerp, Belgium, 2007.
- [21] T. T. Wang and T. F. Quatieri, "Towards Co-channel Separation by 2-D Demodulation of Spectrograms," presented at IEEE Workshops on Applications of
20 Signal Processing to Audio and Acoustics, 2009, *to be presented*.
- [22] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, 2nd ed. New York, NY: John Wiley and Sons, 2001.
- [23] L. Rabiner, M. Cheng, A. Rosenberg, and A. McGonegal, "A Comparative Study of Several Pitch Detection Algorithms," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 23, 1976.
25

[24] Y. Bar-Shalom, X. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation: Theory, Algorithms, and Software*. New York: Wiley, 2001.

- 5 [25] T. Wang and T.F. Quatieri, "High-pitch Formant Estimation by Exploiting Temporal Change of Pitch," IEEE TASL, v18-1, January 2010.

[26] J. Proakis, *Digital Communications*. New York, NY: McGraw-Hill, Inc. 2001.

CLAIMS

What is claimed is:

1. A method for processing an acoustic signal, the method comprising:
preparing a first frequency-related representation of the acoustic
5 signal over time;
computing a two-dimensional transform of plural two-dimensional
localized regions of the first frequency-related representation, each less than
an entire frequency range of the first frequency related representation, to
provide a two-dimensional compressed frequency-related representation with
10 respect to each two dimensional localized region;
identifying at least one pitch for each of the plural regions; and
processing the pitch from the plural regions to provide multiple pitch
estimates over time.
2. The method of Claim 1 wherein the acoustic signal includes at least one
15 audio source signal.
3. The method of Claim 1 wherein the acoustic signal includes a plurality of
audio source signals from independent speech signal sources.
4. The method of Claim 1 wherein the acoustic signal includes voiced and
unvoiced acoustic signals.
- 20 5. The method of Claim 1 wherein the two-dimensional transform includes a
two-dimensional Fourier Transform.
6. The method of Claim 1 further including identifying pitch dynamics
corresponding to the at least one pitch for each of the plural regions and
providing the multiple pitch estimates as a function of processing the pitch
25 and the pitch dynamics.
7. The method of Claim 1 further including representing the at least one pitch
as a function of a vertical distance of a representation of the at least one pitch
from an origin of a frequency-related region.

8. The method of Claim 7 wherein the representation of the at least one pitch includes an impulse train representation.
9. The method of Claim 7 wherein the frequency related region is a Grating Compression Transform region.
- 5 10. The method Claim 1 further including representing the at least one pitch as a function of a vertical distance and a radial angle of a representation of the at least one pitch from an origin of a frequency-related region.
11. The method of Claim 1 further including removing near DC components of the two-dimensional compressed frequency-related representation.
- 10 12. The method of Claim 1 wherein the plural two-dimensional localized regions cover various frequency ranges corresponding to a common time point.
13. The method of Claim 1 wherein the plural two-dimensional localized regions overlap in time.
14. A method for processing a mixed acoustic signal comprised of multiple
15 signals, the method comprising:
 - localizing multiple time-frequency regions of a spectrogram of the mixed acoustic signal to obtain one or more acoustic properties;
 - providing at least one pitch estimate for each of the multiple acoustic signals as a function of combining the acoustic properties; and
 - 20 recovering at least one of the multiple signals as a function of the separate pitch estimate.
15. The method of Claim 14 wherein the one or more acoustic properties include pitch candidates.
16. The method of Claim 14 further including identifying at least one pitch for
25 each localized time-frequency region of the spectrogram.

17. The method of Claim 14 further including demodulating individual signal contents using pitch information to recover information corresponding to an individual acoustic signal.
18. The method of Claim 17 further including demodulating using a series-based sinusoidal demodulation.
19. The method of Claim 18 further including demodulating using one or more sets of individual sinusoidal functions.
20. The method of Claim 18 further including demodulating using one or more sets of sinusoidal series.
21. The method of Claim 17 further including combining the recovered information of the localized regions and reconstructing the at least one of the multiple signals as a function of the combined information.
22. The method of Claim 14 further including estimating model parameters for representing at least one of the multiple signals and recovering the at least one of the multiple signals as a function of the estimated model parameters.
23. The method of Claim 14 wherein the multiple signals include at least one of unvoiced signals, periodic signals, non-periodic signals, and quasi-periodic signals.
24. The method of Claim 14 wherein the multiple signals include two or more voiced signals.
25. The method of Claim 14 wherein the multiple signals include one or more unvoiced signal and a noise signal.
26. The method of Claim 14 wherein the multiple signals include one or more voiced signal and at least one noise signal.
27. The method of Claim 14 wherein the multiple signals include one or more voiced signal and at least one unvoiced signal.

28. The method of Claim 27 further including recovering at least one voiced signal and one unvoiced signal as a function of the separate pitch estimate.
29. The method of Claim 27 further including detecting voiced, unvoiced, or silent time-frequency regions and providing the at least one pitch estimate in
5 an event a voiced time-frequency region is detected.
30. The method of Claim 14 wherein the multiple time-frequency regions include predetermined sizes.
31. The method of Claim 14 wherein the multiple time-frequency regions include variable sizes.
- 10 32. The method of Claim 14 wherein the one or more acoustic properties include an impulse train representation.
33. An apparatus for processing an acoustic signal, the apparatus comprising:
a frequency-related representation preparer that prepares a first frequency-related representation of the acoustic signal over time;
15 a two-dimensional transformer that computes a two-dimensional transform of plural two-dimensional localized regions of the first frequency-related representation, each less than an entire frequency range of the first frequency related representation, to provide a two-dimensional compressed frequency-related representation with respect to each two dimensional
20 localized region;
an identifier that identifies at least one pitch for each of the plural regions; and
a processor that processes the pitch from the plural regions to provide multiple pitch estimates over time.
- 25 34. The apparatus of Claim 33 wherein the acoustic signal includes at least one audio source signal.
35. The apparatus of Claim 33 wherein the audio source signal includes a plurality of audio source signals from independent speech signal sources.

36. The apparatus of Claim 33 wherein the acoustic signal includes voiced and unvoiced acoustic signals.
37. The apparatus of Claim 33 wherein the two-dimensional transform includes a two-dimensional Fourier Transform.
- 5 38. The apparatus of Claim 33 wherein the identifier identifies pitch dynamics corresponding to the at least one pitch for each of the plural regions and the processor processes the multiple pitch estimates as a function of processing the pitch and the pitch dynamics.
- 10 39. The apparatus of Claim 33 wherein the identifier represents the at least one pitch as a function of a vertical distance of a representation of the at least one pitch from an origin of a frequency-related region.
40. The apparatus of Claim 39 wherein the representation of the at least one pitch includes an impulse train representation.
- 15 41. The apparatus of Claim 39 wherein the frequency related region is a Grating Compression Transform region.
42. The apparatus Claim 33 wherein the identifier represents the at least one pitch as a function of a vertical distance and a radial angle of a representation of the at least one pitch from an origin of a frequency-related region.
- 20 43. The apparatus of Claim 33 further including a near DC component remover that removes near DC components of the two-dimensional compressed frequency-related representation.
44. An apparatus for processing a mixed acoustic signal comprised of multiple signals, the apparatus comprising:
a localizer that localizes multiple time-frequency regions of a
25 spectrogram of the mixed acoustic signal to obtain one or more acoustic properties;
a pitch estimate provider that provides at least one pitch estimate for

each of the multiple acoustic signals as a function of combining the acoustic properties; and


a signal recoverer that recovers at least one of the multiple signals as a function of the separate pitch estimate.

- 5 45. The apparatus of Claim 44 wherein the one or more acoustic properties include pitch candidates.
46. The apparatus of Claim 44 further including a pitch identifier that identifies at least one pitch for each localized time-frequency region of the spectrogram.
- 10 47. The apparatus of Claim 44 further including a demodulator that demodulates individual signal contents using pitch information to recover information corresponding to an individual acoustic signal.
48. The apparatus of Claim 47 wherein the demodulator is a series-based sinusoidal demodulator.
- 15 49. The apparatus of Claim 48 wherein the demodulator employs one or more sets of individual sinusoidal functions.
50. The apparatus of Claim 48 wherein the demodulator employs one or more sets of sinusoidal series.
51. The apparatus of Claim 47 further including a combiner that combines the
20 recovered information of the localized regions and reconstructs the at least one of the multiple signals as a function of the combined information.
52. The apparatus of Claim 47 wherein the signal recoverer recovers the at least one of the multiple signals as a function of estimated model parameters that represent the at least one of the multiple signals.


53. The apparatus of Claim 47 wherein the multiple signals include at least one of unvoiced signals, periodic signals, non-periodic signals, and quasi-periodic signals.
54. The apparatus of Claim 47 wherein the multiple signals include two or more
5 voiced signals.
55. The apparatus of Claim 47 wherein the multiple signals include one or more unvoiced signal and a noise signal.
56. The apparatus of Claim 47 wherein the multiple signals include one or more voiced signal and at least one noise signal.
- 10 57. The apparatus of Claim 47 wherein the multiple signals include one or more voiced signal and at least one unvoiced signal.
58. The apparatus of Claim 57 wherein the recoverer recovers at least one voiced signal and one unvoiced signal as a function of the separate pitch estimate.
59. The apparatus of Claim 57 further including a voicing state detector that
15 detects voiced, unvoiced, or silent time-frequency regions and wherein the pitch estimate provider provides the at least one pitch estimate in an event a voiced time-frequency region is detected.
60. The apparatus of Claim 44 wherein the multiple time-frequency regions include variable sizes.
- 20 61. The apparatus of Claim 44 wherein the one or more acoustic properties include an impulse train representation.
62. A method for processing a mixed acoustic signal comprised of multiple signals, the method comprising:
 localizing multiple time-frequency regions of a spectrogram of the
25 mixed acoustic signal to obtain one or more acoustic properties; and

recovering at least one of the multiple signals as a function of the one or more acoustic properties.

FIG. 1A Prior Art



2/42



Prior Art

FIG. 1B




FIG. 1C(b)




FIG. 1C(a)

FIG. 1C(d)





FIG. 1C(c)



$0.5e^{-j\omega} A(\omega + \omega_s \sin \theta, \Omega - \omega_s \cos \theta)$




FIG. 2B





FIG. 2C

5/42

FIG. 3A




FIG. 3B




FIG. 3C





FIG. 3D



6/42

FIG. 4A




FIG. 4B




FIG. 4C





FIG. 4D



500 ↘

	Single Speaker Filtering	Single Speaker Demod.	Mixed Speaker SBSS	Mixed Speaker Demod.	Mixed Speaker TruePhase
SNR (dB)	11.24	12.51	3.62	4.09	5.96

FIG. 5

FIG. 6A




FIG. 6B




FIG. 6C




FIG. 6D




FIG. 7A




FIG. 7B




FIG. 7D






FIG. 7C





11/42




FIG. 9

FIG. 10A





FIG. 10C




FIG. 10B




FIG. 10D




FIG. 11A




FIG. 11B




FIG. 11C




FIG. 11D




FIG. 12A




FIG. 12B




FIG. 12C





FIG. 12D



15/42

	Syn1	Syn2	Syn3	Syn4	Real1	Real2
oracle	0.00	0.00	0.03	0.00	0.00	0.00
clustering	1.74	1.44	1.00	1.08	5.46	7.91
single	5.98	14.97	13.24	9.13	42.18	20.13

Average percent errors across time for each mixture and method

FIG. 13




FIG. 14A




FIG. 14B

FIG. 15A




FIG. 15B




FIG. 15C





FIG. 15D



2-D Fourier Transform

FIG. 16A




FIG. 16B




FIG. 17A




FIG. 17C




FIG. 17B




FIG. 17D




FIG. 18A




FIG. 18C




FIG. 18B






FIG. 18D



22/42

1900



23/42




FIG. 21A




FIG. 21B




FIG. 21C




FIG. 21D




FIG. 22A




FIG. 22B




FIG. 23A





FIG. 23B



27/42

2400

	Avg. E_gross <u>2410</u>	Avg. E_fine <u>2420</u>	Avg. E_total <u>2430</u>
Pruned <u>2440</u>	14.83	6.70	21.53
Raw <u>2450</u>	13.26	6.45	19.71

FIG. 24

2500

	Avg. E_gross <u>2510</u>	Avg. E_fine <u>2520</u>	Avg. E_total <u>2530</u>
Pruned <u>2540</u>	5.58	7.04	12.62
Raw <u>2550</u>	13.29	7.74	21.03

FIG. 25

2600




FIG. 26A

28/42




FIG. 26B




FIG. 26C

29/42




FIG. 26D




FIG. 26E

FIG. 26F

30/42




FIG. 26G




FIG. 26H





FIG. 26I



31/42




FIG. 26J




FIG. 26K

FIG. 26L




FIG. 26M




FIG. 26N





FIG. 26O



33/42




FIG. 27A




FIG. 27B

34/42




FIG. 27C




FIG. 27D

Size/Signal	Synth1	Synth2	Synth3	Synth4
20 ms by 625 Hz	6.95	4.66	7.40	16.37
42 ms by 625 Hz	6.60	4.76	6.02	3.86
20 ms by 687.5 Hz	6.60	4.55	7.45	16.78
20 ms by 812.5 Hz	6.58	4.31	6.91	18.86

FIG. 28




FIG. 29A

FIG. 29B

FIG. 30A




FIG. 30B




FIG. 31A




FIG. 31B







FIG. 32A








FIG. 33




FIG. 34