



Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Appl. Comput. Harmon. Anal. 18 (2005) 186–197

**Applied and
Computational
Harmonic Analysis**

www.elsevier.com/locate/acha

Letter to the Editor

Using edge information in time–frequency representations for chirp parameter estimation

Nicholas N. Bennett ^{a,*}, Naoki Saito ^b

^a Schlumberger-Doll Research, 36 Old Quarry Road, Ridgefield, CT 06877, USA

^b Department of Mathematics, University of California, Davis, One Shields Avenue Davis, CA 95616-8633, USA

Received 14 October 2003; revised 23 June 2004; accepted 2 November 2004

Available online 22 December 2004

Communicated by M.V. Wickerhauser

Abstract

Time–frequency representations of a signal can provide a useful means for obtaining parameter estimates for signals consisting of various chirps. We demonstrate the utility of including edge information extracted from these time–frequency representations when using a Hough transformation to perform this task. In particular, we show that using the edge information: (1) reduces the variance of the chirp parameter estimates in the case where the chirp signal has a single component; and (2) reduces the amount of spurious cross talk results when the signal has multiple chirp components. We further demonstrate a variation of our technique that detects the onset and duration of individual chirp components. We propose this technique as a fast preprocessing step for other algorithms such as maximum likelihood estimation which can provide very accurate parameter estimates.

© 2004 Elsevier Inc. All rights reserved.

Keywords: Local Fourier basis; Hough transform; Edge detection; Uncertainty; Randomized Hough transform; Time–frequency analysis

* Corresponding author.

E-mail addresses: nmbennett@slb.com (N.N. Bennett), saito@math.ucdavis.edu (N. Saito).

URL: <http://www.math.ucdavis.edu/~saito>.

1. Introduction

There is a long tradition of using sonogram-based tools such as the short time Fourier transform and the Wigner–Ville distribution to obtain a graph of the phase of a signal’s components [3,4,12,19]. The use of time–frequency representations such as the local trigonometric, wavelet packet, and adaptive Gabor transforms has drawn further attention to the problem of estimating the instantaneous frequency of a signal or its components [7]. In this paper, we are interested in those problems where there is a parametric model for the phase of the signal’s components and where one would like to obtain estimates of these parameters. Many useful techniques may be found in the literature [1,4,13,17].

The technique we propose is the use of a Hough transformation to process the edges which are extracted from the signal’s local Fourier representation. We make use of a variant of the Hough transform called the randomized Hough transform (RHT) [20] to randomly sample the well-resolved edges in the time–frequency image and to convert these edge groups to parameter estimates. The edge information does carry substantial uncertainty, but we show that its use: (1) reduces the variance of the chirp parameter estimates in the case where the chirp signal has a single component; and (2) reduces the amount of spurious cross talk results when the signal has multiple chirp components. We also propose a second algorithm that uses the uncertainty of the edge information to group edges with consistent orientations.

We do not attempt a direct comparison of our methods with the very large number of other well-known techniques, but instead have decided to provide a means for readers to decide for themselves if our methods are useful for their specific problems by making available the codes required to reproduce the figures in this paper. The codes make use of the Donoho’s WaveLab software and can be found at Naoki Saito’s web page <http://www.math.ucdavis.edu/~saito/software/>. We hope that following WaveLab’s ‘reproducible research’ philosophy [2] aids the progress of this long line of research results recorded in the literature.

2. Problem formulation

We consider discretized signals of the form

$$f(t_l) = \sum_{k=1}^K e^{2\pi i \frac{N}{2} p_k(t_l)} + n_l, \tag{1}$$

where K is the number of chirps in the signal, the t_l is an N point uniform discretization of $[0, 1]$, and n_l are samples from a Gaussian distribution. We shall modify the variance of this Gaussian distribution later to control the signal-to-noise ratio for our problem. The functions p_k we shall consider belong to one of the following classes:

- Linear chirps

$$p_k(t) = \alpha_k t + \tan(\beta_k) \frac{t^2}{2}; \tag{2}$$

- Logarithmic chirps [8]

$$p_k(t) = (\alpha_k - \beta_k)t + \beta_k \log(t + 1); \tag{3}$$

- Sinusoidal chirps [6,11]

$$p_k(t) = \frac{-\alpha_k}{2\pi\beta_k} \cos(2\pi(\beta_k t - \gamma_k)). \quad (4)$$

We shall assume for simplicity that we know a priori to which class the signal's components belong (e.g., we know that we are listening for helicopters or bats, etc.). We shall remove the implicit assumption that the chirp components' duration is over all of $[0, 1]$ when we present Algorithm II. Our initial objective is to determine the number of components K and to estimate the P parameters describing each chirp.

3. Time–frequency representation

Our time–frequency representation shall be constructed using a \sqrt{N} local Fourier basis. This local Fourier basis is constructed from the \sqrt{N} local sine and cosine bases (see [5] for details): for $1 \leq j \leq \sqrt{N}$ and $1 \leq k < \sqrt{N}$ define

$$C_{jk}(t) = b(\sqrt{N}t - j) \cos\left((2k - 1)\frac{\pi}{2}(\sqrt{N}t - j)\right),$$

$$S_{jk}(t) = b(\sqrt{N}t - j) \sin\left((2k - 1)\frac{\pi}{2}(\sqrt{N}t - j)\right),$$

where

$$b(x) = \begin{cases} \cos((\pi/2)t), & \text{if } |t| < 1, \\ 0, & \text{otherwise.} \end{cases}$$

We use these collections and the fast algorithms associated with them to compute a \sqrt{N} local Fourier representation for our signal,

$$\lambda_{j,\pm k} = \int f(t)(C_{jk}(t) \pm iS_{jk}(t)) dt. \quad (5)$$

The time–frequency representation used here is an image with pixel dimensions $2\sqrt{N}$ by \sqrt{N} . To the pixel with index (j, k) , we shall assign intensity $|\lambda_{j,k}|$ and a position in the rectangle $[-1, 1] \times [0, 1]$ given by $(2j - 1, 2k - 1)/\sqrt{N}$. For the various classes of chirps in Eqs. (2) through (4), it is easy to see that the large coefficients of f will appear in the image along ridges which follow the curves $\xi = p'_k(t)$. Though one could apply the variants of the RHT we present below to other time–frequency representations such as the Wigner–Ville distribution, we prefer the rather coarse local Fourier representation so that

- (1) we can handle signals with many samples (e.g., several thousand),
- (2) denoised versions of our signal can be obtained from subsamples of the pixels, and
- (3) we have square pixels in our time–frequency representation.

Other choices would be quite appropriate given additional prior information concerning, for example, the frequency band, dynamic range, or duration of the signal components.

In Fig. 1, the time–frequency representation of signals consisting of the linear chirps $\alpha_1 = 0.1$, $\beta_1 = 0.4$ and $\alpha_2 = -0.1$, $\beta_2 = 0.5$ is shown for $N = 256$.

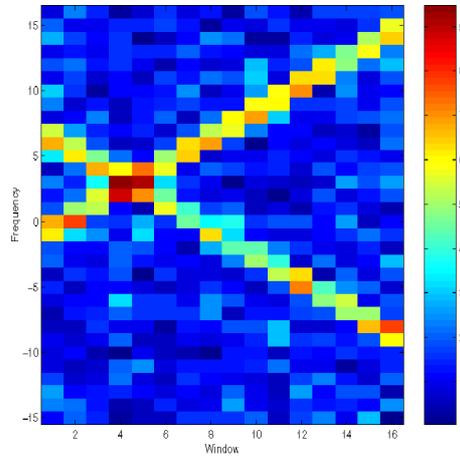


Fig. 1. Time–frequency image derived from the local Fourier representation of a signal having two chirp components with parameters $\alpha_1 = -0.1$, $\beta_1 = \pi/4$, $\alpha_2 = 0.4$, $\beta_2 = -\pi/4$, and $N = 256$.

4. Edge detection and uncertainty

The problem of parameter estimation for our signal has thus been converted to the problem of detecting the parabolas, lines, or other curves $\xi = p'_k(t)$ present in the time–frequency representation.

To detect the edges in our image, we employ the matched filter technique of Nevatia and Babu [18] and the suggested modifications of this technique due to Lyvers and Mitchell [14]. The procedure is basically as follows: at each pixel in the image, compute averages of the pixel values in various directions; find the maximum of these responses and call the corresponding direction the *edge angle*; call the value of this maximum response the *edge magnitude*. In this paper we chose to employ matched filters for the following 8 angles: $\pi/2, \pm 3\pi/8, \pm \pi/4, \pm \pi/8, 0$. In Fig. 2 we show edge angle and magnitude images associated with the time–frequency image displayed in Fig. 1.

We prefer these matched filtering techniques to gradient methods [15] due to the fact that gradient methods are primarily used for detecting step edges (usually the more prevalent type in typical images), whereas we are interested in detecting line edges. The use of step edge information leads to detection of two parameter curves per signal component—one for each side of the ridge of large local Fourier coordinates. Note that due to the size of our image, all of the edge detection computations cost order N computations.

The edge data comes in the form (t_i, ξ_i, θ_i) , where (t_i, ξ_i) is the position of the edge in the time–frequency image and $\tan(\theta_i)$ is the slope of the edge. If this edge falls along the ridge of large coefficients, this data will approximately satisfy the equations

$$\xi_i = p'(t_i), \tag{6}$$

$$\tan(\theta) = p''(t_i). \tag{7}$$

The uncertainty in the positions of the detected edges can arise due to noise in the recording or irregularities in the chirps themselves. Further discussion of the limitations of Eqs. (6) and (7) can be found in [9]. The primary uncertainty in the edge angle information arises from the quantization of the angle

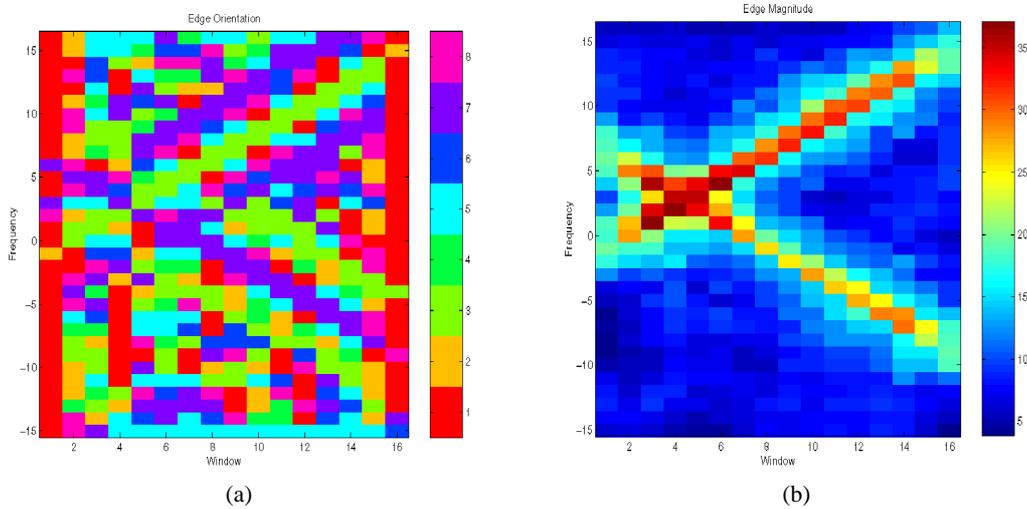


Fig. 2. The edge orientation and magnitude images associated with the time–frequency image displayed in Fig. 1. The edge orientation colormap associates colormap values with edge orientation angles as follows: $1 \leftrightarrow \pi/2$, $2 \leftrightarrow 3\pi/8$, $3 \leftrightarrow \pi/4$, $4 \leftrightarrow \pi/8$, \dots , $8 \leftrightarrow -3\pi/8$.

data according to the number of matched filters employed and is the primary object of study here. For examples discussed in this paper, we considered the edge angle uncertainty to be $\pm\pi/16$.

5. Hough processing

The Hough transform is a well-known method for detecting curves in an image [10]. The particular variant we shall employ here is the randomized Hough transform (RHT) [20]. For a chirp consisting of P unknown parameters, the RHT algorithm selects P pixels at random (from among those pixels with intensity larger than some threshold) and maps them to a single bin in a histogram, usually called an accumulator array, describing a region of the chirp parameter space. Typically, the bin size of the histogram can be estimated from performing a statistical analysis such as we perform in Experiment B below.

We shall make use of the edge information to eliminate spurious votes from the standard RHT procedure as follows. The randomly chosen P pixels in the time–frequency image determine a unique curve in the image. We shall allow the parameter space vote only if the edge information at the chosen pixels is consistent with the derivative of the curve.

5.1. Algorithm I outline

Inputs:

- Samples of signal in equation with P parameters given in Eq. (1);
- Edge magnitude threshold $T_e = 75\%$ of the maximum magnitude;
- Parameter histogram threshold $T_h = 75\%$ of the maximum vote count.

Processing:

- (1) Compute local Fourier coefficients as in Eq. (5) [cost: $\mathcal{O}(N \log N)$].
- (2) Compute edge angle and magnitude information from the local Fourier coefficient image [cost: $\mathcal{O}(N)$].
- (3) Threshold the edge magnitude file (choose edges with intensity $> T_e$).
- (4) Compute the Hough transform of this edge collection. Repeat the following for some user defined number of times:
 - (a) Select at random groups of P edges;
 - (b) Compute a parametric curve passing through the P edge locations;
 - (c) Compute the derivative of this parametric curve at the edge locations;
 - (d) If the edge data is consistent with the derivative information (e.g., if their difference is less than the uncertainty of the edge data), then vote for the curve parameters in the parameter histogram. [Cost: proportional to number of groups processed.]
- (5) Threshold (choose local maxima with score $> T_h$) and rank the local maxima within the accumulator array according to their score [cost: inversely proportional to the histogram bin size].

Output: List of detected phase function parameters.

5.2. Experiments using Algorithm 1

In Experiment A, we ran the Hough processing with and without the edge data for a signal having two linear chirps $\alpha_1 = 0.2$, $\beta_1 = -\pi/6$ and $\alpha_2 = -0.2$, $\beta_2 = \pi/6$. The signal-to-noise ratio (SNR) was 0. We employ WaveLab's definition of SNR,

$$\text{SNR}(f, n) = 20 \log_{10} \left(\frac{\|f\|_2}{\|f - n\|_2} \right), \quad (8)$$

where f is the reference signal and n the noisy signal. In Fig. 3, we show the two accumulator arrays. Note that a common difficulty in using the Hough transform when multiple objects are present in the image is the tendency for spurious results due to cross talk between features of those objects. This is clearly visible in the left subfigure. In our experience, the consistency test required of the edges reduces this spurious cross talk effect. In the software contributed with this paper we show that when choosing $\alpha_1 = 0.1$, $\beta_1 = -\pi/8$ and $\alpha_2 = 0.0$, $\beta_2 = -\pi/16$, the spurious cross talk effect is sufficient to generate an erroneous maximum peak when not using the edge information.

In Experiment B, we characterize the variance of the chirp parameter estimates with and without the use of the edge information. One means of doing this is to compute the mean and standard deviation of parameter votes that are cast in the accumulator array during the Hough processing. The standard deviation gives a measure of how tightly the distribution of votes are around the estimated parameter value. In order to make a more general statement, we compute the average standard deviation of the parameter votes for a large number of experiments that estimate the parameters of signals consisting of a single linear chirp. We segmented the experiments according to the true slope of these linear chirp signals. We display the results of the following processing in Fig. 4. For each choice of $\tan(\beta)$ we computed 1000 signals consisting of a single chirp component with slope $\tan(\beta)$ and random y intercept α with SNR-3. For each signal we computed the mean and standard deviation of the individual parameter space

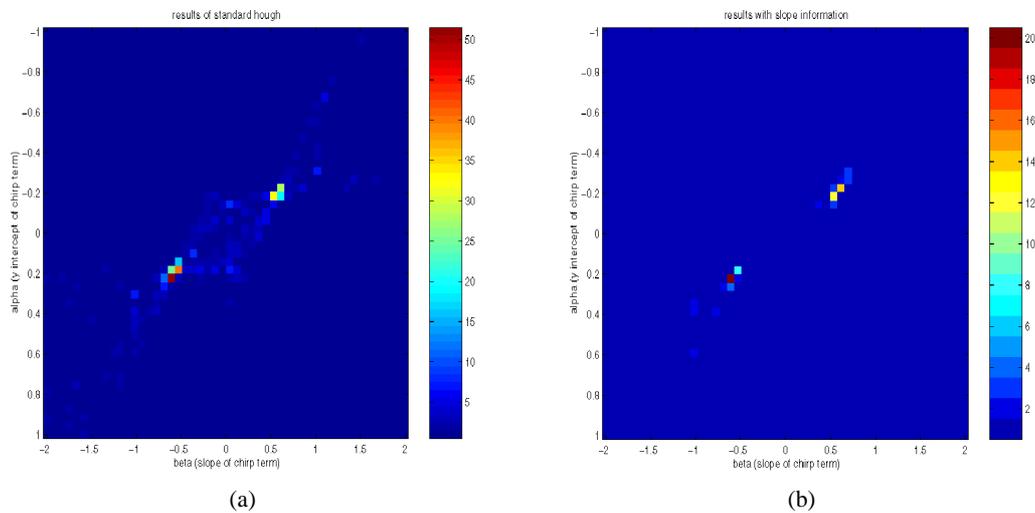


Fig. 3. Experiment A: Images of accumulator arrays for Hough processing with and without the edge information. Note that the ‘cross talk’ between the two chirps in the accumulator array (the two symmetric ridges in the accumulator array running between the two local maxima) has been reduced when making use of the edge information.

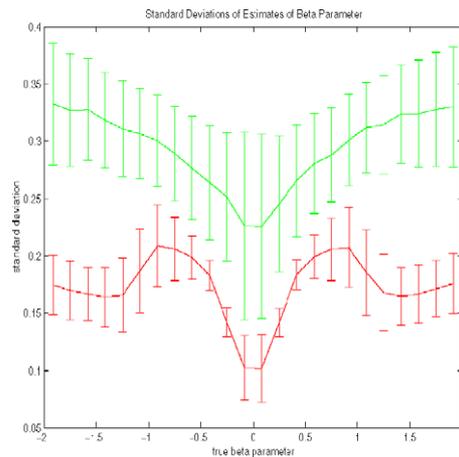


Fig. 4. Experiment B: Comparison of parameter estimate standard deviations with and without the edge information for different linear chirps. The error bars represent the middle two quartiles of the distribution of standard deviations.

votes. In Fig. 4 we display the mean and standard deviation of these standard deviations: the error bars represent the two middle quartiles of the distribution of standard deviations. Note that the use of the edge information decrease the average standard deviations of the parameter estimates, usually by a factor of at least 2. Note that the edge information makes less of a contribution for choices of β which fall between the angles of our matched filter. We use the average standard deviation to determine the bin sizes for our parameter space histograms: typically we use bin sizes which correspond to one half or one quarter of the standard deviation. In the software contributed with this paper, we repeat Experiment B for logarithmic chirps.

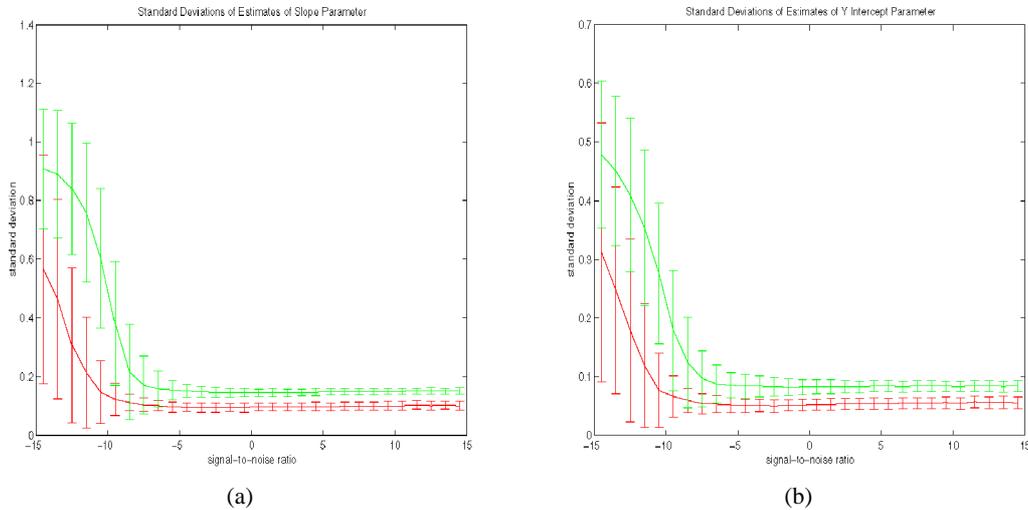


Fig. 5. Experiment C: Comparison of parameter estimate standard deviations with and without the edge information for different SNRs. The error bars represent the middle two quartiles of the distribution of standard deviations.

In Experiment C, we characterize the standard deviation of our linear chirp parameter estimates at different SNRs. We selected $\alpha = -0.2$ and $\beta = \pi/6$ and prepared groups of 500 signals with this chirp and a common SNR. We ran our Hough processing with and without the edge data. For each signal we computed the mean and standard deviation of the individual parameter space votes. In Fig. 5 we display the mean and standard deviation of these standard deviations: as in Fig. 4, the error bars represent the two middle quartiles of the distribution of standard deviations. We observe an increase in the accuracy of our estimation results with increasing SNR until they reach a plateau around SNR 2. This plateau in accuracy is expected because each parameter estimate vote uses only two pixels. These two pixels have a position uncertainty of $\pm 1/\sqrt{N}$, so the parameter estimates will have expected standard deviation which is bounded from below.

6. Using uncertainty to group edges

The uncertainty associated with the edges extracted from the time–frequency representation of a signal can be used to group edges with consistent orientations and, following some of the ideas of [16], lead to the following “greedy” variant of our RHT algorithm.

Instead of immediately computing and voting for candidate curve parameter values whenever we find P consistent edges in our time–frequency image, we shall also test the other edge pixels along the candidate parameter curve in the time–frequency image. We then remove the bright edge pixels (i.e. those with large magnitude) along the parameter curve from the pool of bright edges and continue the process until the number of remaining bright edges gets small or if we have processed too many inconsistent edge pairs.

Using the bright edges along this candidate curve, we can:

- (1) Report the estimated curve parameters (which we can estimate using a least squares fit);

- (2) Compute and report the minimum and maximum time position of these consistent edge pixels and thereby track the onset and termination of the chirp component. This allows the algorithm to be useful in somewhat less tailored settings than those we have described earlier in the paper;
- (3) The thresholds of the RHT variant we present below can be phrased in terms of the phase plane parameter curve characteristics instead in terms of a histogram threshold. For instance, we can ask to detect only chirps of a certain duration;
- (4) We avoid the use of a parameter space histogram whose computational burdens can be substantial when $P > 2$.

6.1. Algorithm II outline

Inputs:

- Samples of signal in Eq. (1) with P parameters;
- Edge magnitude threshold $T_e = 50\%$ of the maximum magnitude;
- Minimum length of candidate curve segment, $T_l = (1/2)\sqrt{N}$;
- Minimum percentage of edge pairs along a candidate curve segment which are consistent, $T_c = 75\%$;
- Maximum number of inconsistent edge pairs N_f .

Processing:

- (1) Compute local Fourier coefficients as in Eq. (5) [cost: $\mathcal{O}(N \log N)$].
- (2) Compute edge angle and magnitude information from the local Fourier coefficient image [cost: $\mathcal{O}(N)$].
- (3) Threshold the edge magnitude file (choose edges with intensity T_e). Call this collection of edges V .
- (4) Compute the Hough transform of this edge collection. Repeat the following until the number of edges in the set V becomes less than some number n_v or until there are at least n_f failures at steps 4d, 4f, and 4g:
 - (a) If $n_f \geq N_f$, exit Hough processing loop. Otherwise, select at random a group, G , of P edges from V ;
 - (b) Compute a parametric curve passing through the P edge locations;
 - (c) Compute the derivative of this parametric curve at the edge locations;
 - (d) Test whether the edge data is consistent with the derivative information (e.g., if their difference is less than the uncertainty of the edge data). If yes, proceed to the next step. Otherwise increment n_f and return to step (4a);
 - (e) Count the number, C , of bright edges which fall along the parametric curve. Call this set \tilde{G} ;
 - (f) Test $C > T_l$. If yes, proceed to the next step. Otherwise increment n_f and return to step (4a);
 - (g) Test if $T_c\%$ of these bright edges, \tilde{G} , are consistent with the derivative information. If yes, proceed to the next step. Otherwise increment n_f and return to step (4a);
 - (h) $V = V \setminus \tilde{G}$;
 - (i) Compute a least-squares estimate of curve parameter values using the collection \tilde{G} ;
 - (j) Compute a the minimum and maximum of the window locations for the edges in \tilde{G} .
 [Cost: proportional to number of edge groups processed.]

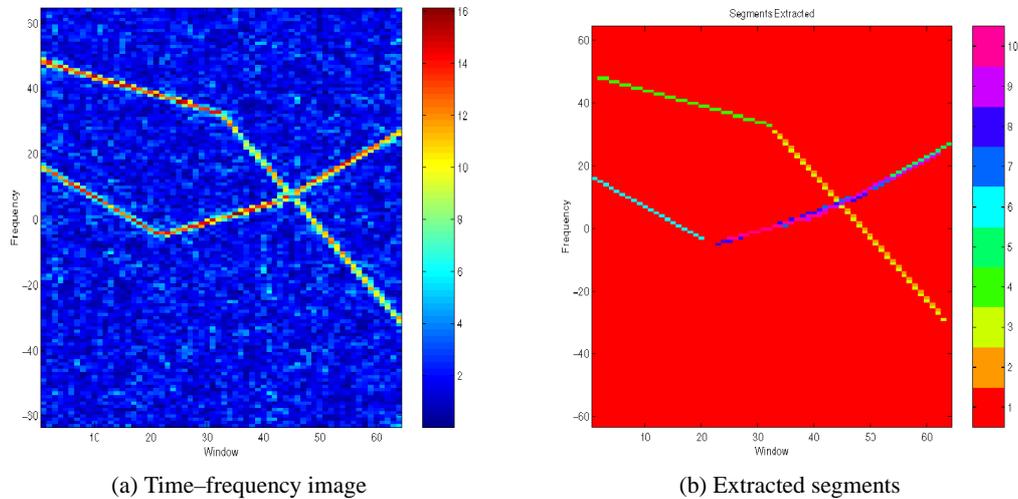


Fig. 6. Experiment E: Time–frequency image of signal consisting of several linear chirps so that their signatures form a series of segments and an image of the extracted segments.

Output: List of detected phase function parameters along with their onset and termination windows.

6.2. Experiments using Algorithm II

In Experiment E, we perform Algorithm II on a group of piecewise linear chirps. In Fig. 6, we show the time–frequency image of our signal that consists of a series of chirp segments and the detected segments. Note that we have multiple returns for some of the chirp segments.

In Experiment F, we perform Algorithm II on a group of sinusoidal chirps. In this case, we must manage a new obstacle in that the equation for the chirps' phase is nonlinear, and given three points in the time–frequency image, there can be more than one sinusoid passing through the three points (in the software supplied with this paper, we illustrate this state of affairs). In step (4b), we introduce an inner loop over these multiple sinusoidal curves and thereby sort through which of them are consistent with the edge information. In Fig. 7a we display the time–frequency image for a signal consisting of two sinusoidal chirp components with SNR-5. In Fig. 7b we show the detected components. Note again that we do have multiple returns for each component.

7. Discussion

In this paper we have presented a collection of new algorithms which employ a local Fourier basis and variants of the randomized Hough transform to compute estimates for parameters of multiparameter chirps. We have made particular emphasis on the usefulness of edge information in lowering the variance of parameter estimates and in decreasing the amount of spurious cross talk observed in the traditional Hough processing techniques. We also shown that the uncertainty of the edge information can be used to group edges with consistent orientations.

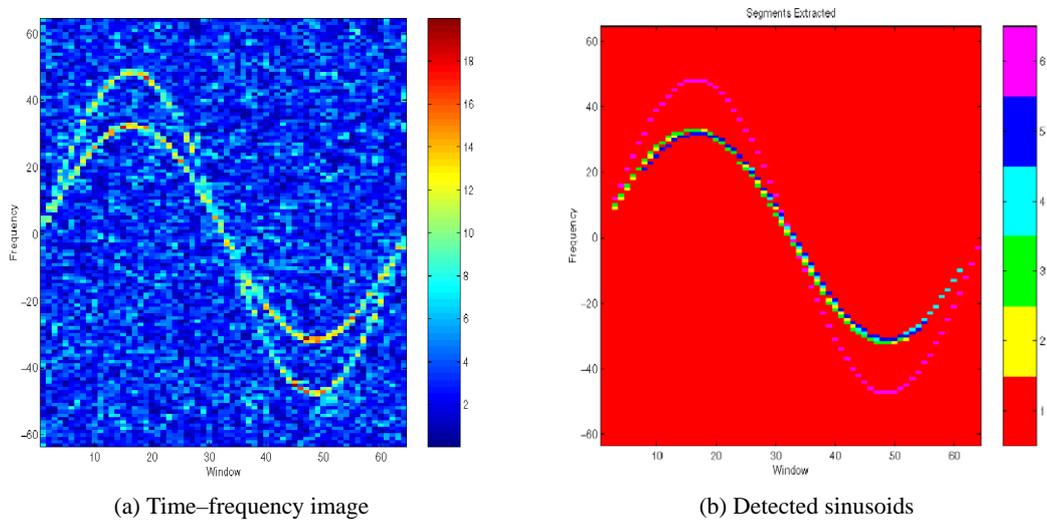


Fig. 7. Experiment F: Time–frequency image of signal consisting of two sinusoidal chirp components with parameters $(\alpha, \beta, \gamma) = (0.5, 1.0, 0.0), (0.75, 1.0, 0.0)$ and SNR-5. The sinusoids detected using Algorithm II.

References

- [1] S. Barbarossa, Analysis of multicomponent LFM signals by a combined Wigner–Hough transform, *IEEE Trans. Signal Process. Lett.* 3 (4) (1996) 1511–1515.
- [2] J. Buckheit, D.L. Donoho, Wavelab and reproducible research, in: A. Antoniadis (Ed.), *Wavelets and Statistics*, Springer-Verlag, New York, 1995.
- [3] R. Carmona, W. Hwang, B. Torresani, *Practical Time–Frequency Analysis*, Academic Press, 1998.
- [4] F. Cohen, S. Kadambe, G. Boudreaux-Bartels, Tracking of unknown nonstationary chirp signals using unsupervised clustering in the Wigner distribution, *IEEE Trans. Signal Process.* 41 (11) (1993) 3085–3101.
- [5] R. Coifman, Y. Meyer, Remarques sur l’analyse de Fourier a fenetre, *C. R. Acad. Sci. Sér. I* 312 (1991) 259–261.
- [6] N. Colin, G.E. Michel, M. Pecot, Wavelet analysis of helicopter radar signals, in: *Progress in Wavelet Analysis*, 1992.
- [7] I. Daubechies, F. Planchon, Adaptive Gabor transforms, *Appl. Comput. Harmon. Anal.* 13 (2002) 1–21.
- [8] B. Escudie, Wavelet analysis of asymptotic signals: a tentative model for Bat sonar receiver, in: *Proceedings of the Conference on Wavelets*, Marseilles, France, 1992.
- [9] P. Flandrin, Time–frequency/time-scale analysis, in: *Wavelet Analysis and Its Applications*, vol. 10, Academic Press, 1998.
- [10] J. Illingworth, J. Kittler, A survey of the Hough transform, *Comput. Vision Graph. Image Process.* 44 (1) (1988) 87–116.
- [11] J. Innocent, B. Torresani, Wavelets and binary coalescences detection, *Appl. Comput. Harmon. Anal.* 4 (1997) 113–116.
- [12] S. Kadambe, R. Orr, Instantaneous frequency estimation using the cross term deleted Wigner distribution, in: *Proceedings of IEEE SP Symposium on Time–Frequency and Time-Scale Analysis*, 1996, pp. 289–292.
- [13] S.M. Kay, *Fundamentals of Statistical Signal Processing*, Prentice Hall, 1993.
- [14] E.P. Lyvers, O.R. Mitchell, Precision edge contrast and orientation estimation, *IEEE Trans. Pattern Anal. Mach. Intelligence* 10 (6) (1988) 927.
- [15] S. Mallat, S. Zhong, Characterization of signals from multiscale edges, *IEEE Trans. Pattern Anal. Mach. Intelligence* 14 (7) (1992) 710–732.
- [16] R. McLaughlin, Randomized Hough transform: improved ellipse detection with comparison, *Pattern Recogn. Lett.* 19 (1998) 299–305.
- [17] M. Morvidone, B. Torresani, Time-scale approach for chirp detection, *Intern. J. Wavelet Multiresolut. Process.*, submitted for publication.

- [18] R. Nevatia, K.R. Babu, Linear feature extraction and description, *Comput. Graph. Image Process.* 13 (3) (1980) 257–269.
- [19] K. Wong, Q. Jin, Estimation of the time-varying frequency of a signal: the Cramer–Rao bound and the application of the Wigner distribution, *IEEE Trans. ASSP* 38 (3) (1990).
- [20] L. Xu, E. Oja, P. Kultanen, A new curve detection method: randomized Hough transform (RHT), *Pattern Recogn. Lett.* 11 (1990) 331–338.