



Time-frequency filtering based on spectrogram zeros

Patrick Flandrin

► To cite this version:

| Patrick Flandrin. Time-frequency filtering based on spectrogram zeros. 2015. ensl-01121522

HAL Id: ensl-01121522

<https://ens-lyon.hal.science/ensl-01121522>

Preprint submitted on 2 Mar 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Time-frequency filtering based on spectrogram zeros

Patrick Flandrin

École Normale Supérieure de Lyon,
Laboratoire de Physique (UMR 5672 CNRS),
46 allée d'Italie, 69364 Lyon Cedex 07 France
`flandrin@ens-lyon.fr`

March 2, 2015

Abstract

For a proper choice of the analysis window, a short-time Fourier transform is known to be completely characterized by its zeros, which coincide with those of the associated spectrogram. A simplified representation of the time-frequency structure of a signal can therefore be given by the Delaunay triangulation attached to spectrogram zeros. In the case of multicomponent AM-FM signals embedded in white Gaussian noise, it turns out that each time-frequency domain attached to a given component can be viewed as the union of adjacent Delaunay triangles whose edge length is an outlier as compared to the distribution in noise-only regions. Identifying such domains offers a new way of disentangling the different components in the time-frequency plane, as well as of reconstructing the corresponding waveforms.

Keywords — time-frequency analysis; filtering; spectrogram; Delaunay triangulation

1 Introduction

Spectrograms—i.e., squared magnitudes of Short-Time Fourier Transforms (STFTs)—are among the simplest and most natural tools for performing a time-frequency analysis of signals [8]. In the case of AM-FM-type signals with a limited number of components, spectrograms are relatively sparse representations, with a few energy ribbons localized along the time-frequency trajectories of the different components¹. Filtering those components is generally achieved by identifying the time-frequency domains defining the ribbons, and then reconstructing waveforms by inverting the transform after masking. This long-standing question has recently received a renewed interest, either because of the development of specific techniques such as synchrosqueezing [1, 6, 22] for which isolating domains of influence is a pre-requisite to reconstruction [20], or because of new proposals such as “contours” [19] for defining basins of attraction attached to components.

In most cases, the rationale for identifying “signal regions” is based on local energy considerations and/or curves such as ridges [7] which are supposed to capture local energy concentration. In contrast with such approaches based on *large values* of the representation, we will propose here to make use of *zeros* as characteristic points.

¹Sparser representations can even be achieved by means of techniques such as reassignment [10, 13], synchrosqueezing [1, 6, 22] or ℓ_1 -minimization [11], but this is not the main point of this paper.

The paper is organized as follows. Section 2 is devoted to STFT/spectrogram, with basics recalled in Section 2.1 and elementary facts on the white Gaussian noise case in Section 2.2. Section 3 then discusses more specifically the role played by zeros in a STFT/spectrogram: Section 3.1 justifies the complete representation they offer for a proper choice of the short-time window, whereas Section 3.2 suggests the Delaunay triangulation based on zeros as a way of getting a simplified description. This paves the way for the new filtering approach that is discussed and illustrated in Section 4.

2 STFT and spectrogram

2.1 Definitions and basics

Given a signal $x(t)$ and a window $h(t)$, the *Short-Time Fourier Transform* (STFT) $F_x^{(h)}(t, \omega)$ is classically defined as the inner product between $x(t)$ and shifted versions (in time and frequency) of $h(t)$:

$$F_x^{(h)}(t, \omega) = \langle x, \mathbf{T}_{t\omega} h \rangle, \quad (1)$$

where $\mathbf{T}_{t\omega}$ stands for some joint time-frequency shift operator. Since individual time and frequency shifts do not commute, this operator is not unique and the transform is defined up to a pure phase term. For a sake of simplicity and symmetry, we will choose here this phase term by identifying $\mathbf{T}_{t\omega}$ with the *Weyl operator* [4, 5]:

$$(\mathbf{T}_{t\omega} h)(s) = h(s - t) \exp \left\{ i\omega \left(s - \frac{t}{2} \right) \right\}, \quad (2)$$

thus ending up with the explicit definition:

$$F_x^{(h)}(t, \omega) = \int_{-\infty}^{+\infty} x(s) \overline{h(s - t)} \exp \left\{ -i\omega \left(s - \frac{t}{2} \right) \right\} ds. \quad (3)$$

The corresponding *spectrogram* (which does not depend upon the choice of the phase term) simply follows as:

$$S_x^{(h)}(t, \omega) = \left| F_x^{(h)}(t, \omega) \right|^2. \quad (4)$$

If both the signal and the window belong to $L^2(\mathbb{R})$ —i.e., are of finite energy—, the (complex-valued) STFT is an isometry from $L^2(\mathbb{R})$ to $L^2(\mathbb{R}^2)$:

$$\iint_{-\infty}^{+\infty} F_x^{(h)}(t, \omega) \overline{F_y^{(h)}(t, \omega)} dt \frac{d\omega}{2\pi} = \|h\|_2^2 \langle x, y \rangle, \quad (5)$$

from which it follows that the associated (real-valued) spectrogram can be seen as an energy distribution in the time-frequency plane, since

$$\frac{1}{\|h\|_2^2} \iint_{-\infty}^{+\infty} S_x^{(h)}(t, \omega) dt \frac{d\omega}{2\pi} = \|x\|_2^2. \quad (6)$$

It is well-known that $F_x^{(h)}(t, \omega)$ and $S_x^{(h)}(t, \omega)$ are not any 2D functions in the sense that, by construction, they inherit some structure from their definition (1) as inner product between the analyzed signal $x(t)$ and the

family of analyzing waveforms that are all derived from the window $h(t)$ by shifts in both time and frequency. More precisely, the STFT (3) satisfies the reproducing identity:

$$F_x^{(h)}(t', \omega') = \int \int_{-\infty}^{+\infty} K(t', \omega'; t, \omega) F_x^{(h)}(t, \omega) dt \frac{d\omega}{2\pi}, \quad (7)$$

in which the reproducing kernel $K(t', \omega'; t, \omega)$ is (up to a complex-valued multiplicative term) nothing but the STFT of the analyzing window:

$$K(t', \omega'; t, \omega) = \frac{1}{\|h\|_2^2} \exp \left\{ \frac{i}{2}(\omega t' - \omega' t) \right\} F_h^{(h)}(t' - t, \omega' - \omega). \quad (8)$$

An equivalent formulation amounts to rewriting $F_h^{(h)}(t, \omega)$ as $A_h(\omega, t)$, with

$$A_x(\xi, \tau) = \int_{-\infty}^{+\infty} x \left(\theta + \frac{\tau}{2} \right) \overline{x \left(\theta - \frac{\tau}{2} \right)} \exp\{-i\xi\theta\} d\theta \quad (9)$$

the so-called *ambiguity function* [8]. Doing so, it is easy to justify that any STFT (or spectrogram) has necessarily some local redundancy since the reproducing kernel (8) cannot be arbitrarily peaked in both time and frequency. This follows from general uncertainty relations attached to ambiguity functions in terms of support and/or volume. In their most elaborate forms, these relations involve L^p norms and lead to inequalities such as [18]

$$\begin{cases} \|A_x\|_p \geq B_p \|x\|_2^2 & \text{for } p < 2 \\ \|A_x\|_p \leq B_p \|x\|_2^2 & \text{for } p > 2 \end{cases} \quad (10)$$

with $B_p = (2/p)^{1/p}$, whereas $\|A_x\|_2 = \|x\|_2^2$. In all cases, inequalities are sharp, with the bound attained by Gaussians. Following Gabor's terminology [12], those functions that correspond to maximally concentrated spectrograms are referred to as "logons".

In contrast with these *volume* inequalities, a complementary form of uncertainty can be obtained regarding *support* [14]. Indeed, it is known that a signal cannot be compactly supported in both time and frequency, which prevents the squared magnitude of an ambiguity function to concentrate all of the signal energy (as in (6), with $h = x$) in a domain Ω of finite area. Therefore, if we remark that applying the Cauchy-Schwarz inequality to (9) guarantees that $|A_x(\xi, \tau)| \leq \|x\|_2^2$, we can conclude that, if

$$\iint_{\Omega} |A_x(\xi, \tau)|^2 d\tau \frac{d\xi}{2\pi} \geq (1 - \epsilon) \|x\|_2^2 \quad (11)$$

for some $\epsilon > 0$, then we have necessarily

$$\text{supp}(\Omega) \geq 1 - \epsilon. \quad (12)$$

The consequence of these inequalities (volume and support) is that the reproducing kernel (8) has necessarily some non-zero extension that controls the local redundancy of the STFT and of the corresponding spectrogram. In the particular case of the (unit energy) Gaussian window²

$$g(t) = \pi^{-1/4} \exp\{-t^2/2\}, \quad (13)$$

which is referred to as "circular" since

$$A_g(\xi, \tau) = \exp \left\{ -\frac{1}{4}(\xi^2 + \tau^2) \right\}, \quad (14)$$

the reproducing kernel is maximally concentrated and defines an "influence domain" which is circular and whose radius is given by some effective area attached to the 2D Gaussian function of variance 2.

²In the Physics literature (see, e.g., [16]), the corresponding spectrogram is referred to as the "Husimi distribution function" [15].

2.2 Spectrogram of white Gaussian noise

Whereas the spectrogram is classically defined for finite energy, deterministic signals, it can also be used for the analysis of finite power, harmonizable random processes [8]³. More specifically, we will restrict to the idealized case of zero-mean, analytic, white Gaussian noise $n(t)$ such that:

$$\mathbb{E}\{n(t)\overline{n(t')}\} = \gamma_0 \delta(t - t') \quad ; \quad \mathbb{E}\{n(t)n(t')\} = 0. \quad (15)$$

It readily follows from (4) and (15) that the expected value of the spectrogram is in this case constant:

$$\mathbb{E}\left\{S_n^{(h)}(t, \omega)\right\} = \gamma_0^2, \quad (16)$$

whereas the covariance between spectrogram values at two different locations in the time-frequency plane only depends on the corresponding lags in both time and frequency, according to the relation:

$$\text{cov}\left\{S_n^{(h)}(t, \omega), S_n^{(h)}(t', \omega')\right\} = \gamma_0^2 S_h^{(h)}(t' - t, \omega - \omega'). \quad (17)$$

In the specific case of the circular Gaussian window (13), this covariance takes on the simple form

$$\text{cov}\left\{S_n^{(g)}(t, \omega), S_n^{(g)}(t', \omega')\right\} = \gamma_0^2 \exp\left\{-\frac{1}{2}d^2((t, \omega), (t', \omega'))\right\} \quad (18)$$

where

$$d((t, \omega), (t', \omega')) = \sqrt{(t - t')^2 + (\omega - \omega')^2} \quad (19)$$

measures the Euclidian distance in the plane between the two considered points.

As a function of this only distance, the spectrogram of white Gaussian noise can then be considered as a second-order homogeneous (or stationary) field. This homogeneity property carries over to characteristic points of the surface (such as extrema, be they local maxima or zeros). However, due to the reproducing kernel structure recalled above, the distribution of those characteristics points is expected to be constrained as well.

3 Spectrogram zeros

3.1 The Bargmann connection

Time and frequency are usually considered either independently or jointly, but it might be interesting to see them as coordinates of a complex-valued variable, thus identifying the time-frequency plane with the complex plane. Doing so by introducing $z = \omega + it$, a direct calculation shows that, when evaluated with the circular Gaussian window $g(t)$ defined in (13), the STFT (3) can be re-written as:

$$F_x^{(g)}(t, \omega) = \exp\left\{-\frac{1}{4}|z|^2\right\} \mathcal{F}_x(z), \quad (20)$$

where

$$\mathcal{F}_x(z) = \int_{-\infty}^{+\infty} A(z, s) x(s) ds \quad (21)$$

³We will here formally apply the definition (4) to finite duration realizations of such stochastic processes, looking at statistical properties of the corresponding spectrogram characteristics.

and

$$A(z, s) := \pi^{-\frac{1}{4}} \exp \left\{ -\frac{1}{2}s^2 - isz + \frac{1}{4}z^2 \right\}. \quad (22)$$

This corresponds to the Bargmann factorization of the STFT, with (21) the Bargmann transform [2], whose kernel is given by (22). One interest of such a companion formulation for the STFT is that (21) is an entire function, with consequences on the structure of the STFT and the associated spectrogram. More specifically, since the circular Gaussian window (13) is normalized so as to be of unit energy, this immediately results in the upper bound $|F_x^{(g)}(t, \omega)| \leq \|x\|$. Together with the factorization (20), this leads to

$$|\mathcal{F}_x(z)| \leq \|x\| \exp \left\{ \frac{1}{4}|z|^2 \right\}, \quad (23)$$

i.e., to the fact that $\mathcal{F}_x(z)$ is an entire function of order 2 [3]. As a consequence, it admits a Weierstrass-Hadamard factorization of the form [3, 13, 16, 23]

$$\mathcal{F}_x(z) = z^m \exp\{C_0 + C_1 z + C_2 z^2\} \prod_{n=1}^{\infty} \left(1 - \frac{z}{z_n}\right) \exp \left\{ \frac{z}{z_n} + \frac{1}{2} \left(\frac{z}{z_n}\right)^2 \right\}, \quad (24)$$

where the variables $z_n = \omega_n + it_n$ stand for the (infinitely many) zeros of the Bargmann transform which, by construction, also correspond to the zeros of the STFT and of the spectrogram. As commented in ([23], Appendix B), the quadratic polynomial that enters the exponential in front of the infinite product allows for simple geometrical transformations such as normalization (C_0), translation/rotation (C_1) and squeezing (C_2), while the integer m is an extra degree of freedom corresponding to a possible m -fold zero at the origin of the plane (this is, e.g., the case for Hermite functions whose order is strictly positive). Although (24) is unlikely to be used as such for a possible reconstruction, its meaning is that the Bargmann transform (and, hence, the associated STFT/spectrogram) is nevertheless completely characterized by the distribution of its zeros.

3.2 Delaunay triangulation

Since zeros completely characterize a STFT, it is natural to consider them as a 2D point process in the time-frequency plane, with distinctive properties attached to the specific nature of the analyzed signal. We can therefore get a simplified, geometrical description of the time-frequency structure of a signal by looking at diagrams connecting STFT/spectrogram zeros—the so-called “stellar representation” in Quantum Mechanics [17] (see also [13] for a related time-frequency perspective)—, the simplest one being the *Delaunay triangulation* [21]⁴. An example of a Delaunay triangulation attached to the collection of STFT/spectrogram zeros in the case of white Gaussian noise is given in Figure 1 (left diagram). Since the stationarity of the analyzed white Gaussian noise results in the homogeneity of the 2D random field defined by the STFT/spectrogram (see Section 2.2), the distribution of zeros is itself homogeneous all over the plane, a situation that is expected to be broken whenever some signal—with a coherent time-frequency structure, such as a frequency modulation—happens to be superimposed. As evidenced in the same Figure 1 (right diagram), this is exactly what happens: when an AM-FM chirp is added to the noise of the left diagram, the noise-only regions remain unaffected whereas the “signal domain” is characterized not only by large spectrogram values but also by Delaunay triangles that are more elongated than in noise-only regions.

⁴Let us recall that a Delaunay triangulation is dual of a Voronoi tessellation in which each cell attached to a given point considered as a center, consists in all points that are closer to this center than to any other one.

4 Time-frequency filtering

4.1 Rationale

Both the theoretical considerations of the previous sections and the evidences of Figure 1 suggest that signal domains can be identified by looking at Delaunay triangles that depart from the expected behavior attached to noise, thus calling for a characterization of this reference situation. In this respect, Figure 2 displays the distribution of edge lengths of Delaunay triangles constructed upon STFT/spectrogram zeros in the nominal case of white Gaussian noise. What is evidenced is that such a length is essentially bounded above by a maximum value $L_{max} \sim 2.16$, with a very low probability to exceed 2 (referring as $|e_{mn}|$ the distance $d(z_m, z_n)$ between any two zeros z_m and z_n , a numerical evaluation shows that $\text{Prob}\{|e_{mn}| > 2\} \sim 10^{-3}$). Selecting Delaunay triangles on the basis of thresholding their maximum edge length is therefore a simple way of identifying elementary signal domains whose concatenation defines supports—delineated by zeros—for time-frequency 1/0 masks to be applied to STFT prior reconstruction of the corresponding signal components.

4.2 Algorithm

Based on the elements obtained above, the time-frequency filtering algorithm is quite straightforward and can be summarized as follows:

1. Perform Delaunay triangulation over STFT zeros z_m ;
2. Identify outlier edges such that $|e_{mn}| = d(z_m, z_n) > \eta = 2$;
3. Keep triangles with at least one outlier edge;
4. Group adjacent such triangles in connected, disjoint domains \mathcal{D}_j ;
5. Multiply STFT with labeled 1/0 masks $\mathbf{1}_{\mathcal{D}_j}(t, \omega)$;
6. Reconstruct the disentangled components, domain by domain.

The last (reconstruction) step can be achieved by using either the standard formula

$$x_j(t) = \iint_{(t, \omega) \in \mathcal{D}_j} F_x^{(h)}(t, \omega) h(s - t) \exp \left\{ i\omega \left(s - \frac{t}{2} \right) \right\} dt \frac{d\omega}{2\pi} \quad (25)$$

or the simplified form

$$x_j(t) = \frac{1}{h(0)} \int_{(t, \omega) \in \mathcal{D}_j} F_x^{(h)}(t, \omega) \frac{d\omega}{2\pi} \quad (26)$$

that is preferred in the synchrosqueezing framework [22].

Some remarks can be made about this approach and how to use it. The first point is that, whatever the level that may be very weak, noise contributions are supposed to actually exist all over the plane. In the case where truly zero values would be encountered in the signal under study (e.g., due to zero-padding introduced so as to limit end-point effects), some very small quantity of noise has to be added. From a different perspective, the outlier selection (Point 2. of the algorithm) relies on some threshold η whose value determines a trade-off between true detection and false positives: choosing $\eta = 2$ for this threshold proved quite effective on the considered examples, but this could be studied further. Incidentally, one can remark that the identified outliers also comprise the very long edges that are at the periphery of the domain and that close the convex hull of all zeros. A simple recipe is not to consider those edges such that $|e_{mn}| \gg L_{max}$, where $L_{max} \sim 2.16$ corresponds to the observed upper bound for edge length (see Figure 2).

4.3 Examples

The first example consists in a Hermite function whose time-frequency “trajectory” is known to be a circle [9]. The overall filtering procedure (triangulation, selection of outliers, grouping and masking) is summarized in Figure 3, with the corresponding reconstruction result in Figure 4.

The second example corresponds to the classical benchmark “bat signal”⁵, whose spectrogram and reassigned spectrogram are plotted in Figure 5. The Delaunay selection of the domains corresponding to the different components is presented in Figure 6, with individual reconstructions of the 3 main ones plotted in Figure 7, together with their recombination to be compared to the complete waveform.

Acknowledgements

This work is part of the ASTRES Project which is supported by Agence Nationale de la Recherche under grant ANR-13-BS03-0002-01.

References

- [1] F. Auger, P. Flandrin, Y.T. Lin, S. McLaughlin, S. Meignen, T. Oberlin and H.T. Wu, “Time-frequency reassignment and synchrosqueezing,” *IEEE Sig. Proc. Mag.*, Vol. 30, No. 6, pp. 32-41, 2013.
- [2] V. Bargmann, “On a Hilbert space of analytic functions and an associated integral transform,” *Commun. Pure Appl. Math.*, Vol. 14, pp. 187–214, 1961.
- [3] R.P. Boas, *Entire Functions*, Academic Press, New York, 1954.
- [4] E. Chassande-Mottin, I. Daubechies, F. Auger, and P. Flandrin, “Differential reassignment,” *IEEE Signal Proc. Lett.*, Vol. 4, No. 10, pp. 293–294, 1997.
- [5] I. Daubechies, *Ten Lectures on Wavelets*, SIAM, Philadelphia (PA), 1992.
- [6] I. Daubechies, J. Lu, and H.-T. Wu, “Synchrosqueezed wavelet transforms: An empirical mode decomposition-like tool,” *Appl. and Comp. Harm. Anal.*, Vol. 30, No. 1, pp. 243–261, 2011.
- [7] N. Delprat, B. Escudié, P. Guillemain, R. Kronland-Martinet, P. Tchamitchian, and B. Torr  sani, “Asymptotic wavelet and Gabor analysis: extraction of instantaneous frequencies,” *IEEE Trans. on Info. Theory*, Vol. 38, No. 2, pp. 644–664, 1992.
- [8] P. Flandrin, *Time-Frequency/Time-Scale Analysis*, Academic Press, 1999.
- [9] P. Flandrin, “A note on reassigned Gabor spectrograms of Hermite functions,” *J. Fourier Anal. Appl.*, Vol. 19, No. 2, pp. 285-295, 2012.
- [10] P. Flandrin, F. Auger, and E. Chassande-Mottin, “Time-frequency reassignment ? From principles to algorithms,” in *Applications in Time-Frequency Signal Processing*, A. Papandreou-Suppappola, Ed., chapter 5, pp. 179–203, CRC Press, Boca Raton (FL), 2003.

⁵Thanks to C. Condon, K. White and A. Feng of the Beckman Institute of the University of Illinois for the bat data and for permission to use it in this paper.

- [11] P. Flandrin, P. Borgnat, “Time-frequency energy distributions meet compressed sensing,” *IEEE Trans. on Sig. Proc.*, Vol. 58, No. 6, pp. 2974–2982, 2010.
- [12] D. Gabor, “Theory of communication,” . *J. IEE*, Vol. 93, pp. 429–457, 1946.
- [13] T. J. Gardner and M. O. Magnasco, “Sparse time-frequency representations,” *Proc. Nat. Acad. Sc.*, Vol. 103, No. 16, pp. 6094–6099, 2006.
- [14] K. Gröchenig, *Foundations of Time-Frequency Analysis*, Birkhäuser, Boston (MA), 2011.
- [15] K. Husimi, “Some formal properties of the density matrix,” *Proc. Phys. Math. Soc. Jpn*, Vol. 22, pp 264–314, 1940.
- [16] H.J. Korsch, C. Muller, and H. Wiescher, “On the zeros of the Husimi distribution,” *J. Phys. A: Math. Gen.*, Vol. 30, pp. L677–L684, 1997.
- [17] P. Leboeuf and A. Voros, “Chaos-revealing multiplicative representation of quantum eigenstates,” *J. Phys. A Math. Gen.*, Vol. 23, pp. 1765–1774, 1990.
- [18] E.H. Lieb, “Integral bounds for radar ambiguity functions and Wigner distributions,” *J. Math. Phys.*, Vol. 31, pp. 594–599, 1990.
- [19] Y. Lim, B.G. Shinn-Cunningham and T.J. Gardner, “Sparse contour representations of sound,” *IEEE Signal Proc. Lett.*, Vol. 19, No. 10, pp. 684–687, 2012.
- [20] S. Meignen, T. Oberlin, and S. McLaughlin, “On the mode synthesis in the synchrosqueezing method,” *Proc. European Signal Proc. Conf. EUSIPCO-12*, Bucharest (R0), 2012.
- [21] A. Okabe, B. Boots, K. Sugihara, and S.N. Chiu, *Spatial Tessellations — Concepts and Applications of Voronoi Diagrams* (2nd ed.), John Wiley, 2000.
- [22] G. Thakur, H.-T. Wu, “Synchrosqueezing-based recovery of instantaneous frequency from nonuniform samples,” *SIAM J. Math. Anal.*, Vol. 43, pp. 2078–2095, 2011.
- [23] M. Toda, “Phase retrieval problem in quantum chaos and its relation to the origin of irreversibility I,” *Physica D*, Vol. 59, pp. 121–141, 1992.




Figure 1: **Delaunay triangulation.** Left diagram: In the case of white Gaussian noise, the Delaunay triangulation (white lines) constructed upon the zeros of the spectrogram (white dots) reveals an homogeneous distribution of random triangles. Right diagram: when a signal is superimposed to noise, the distribution of the Delaunay triangles remains unaffected in noise-only regions, whereas the “signal domain” is characterized by more elongated triangles.

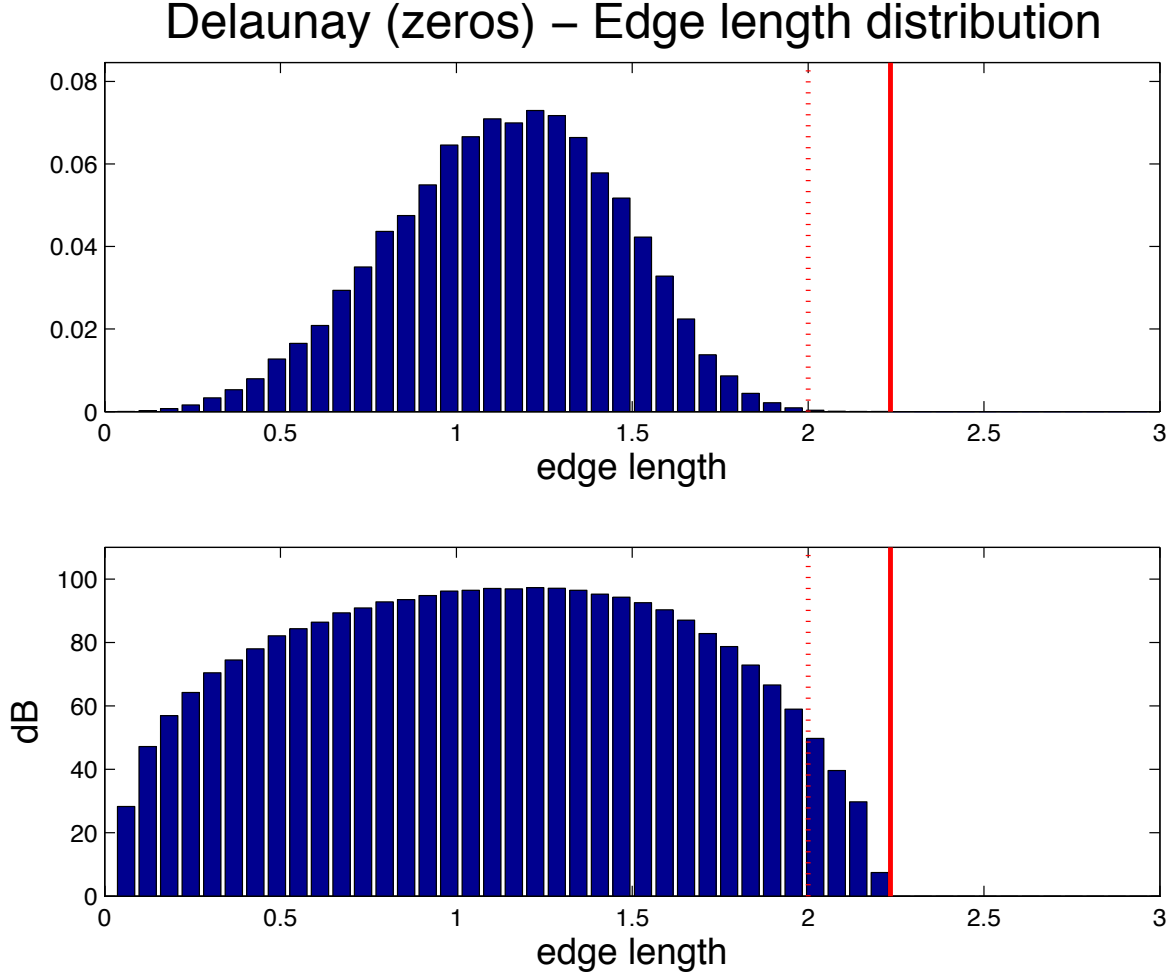


Figure 2: **Delaunay triangulation — Distribution of edge lengths.** In the case of white Gaussian noise, the distribution of edge lengths in Delaunay triangles constructed upon STFT/spectrogram zeros (top diagram: linear scale; bottom diagram: logarithmic scale) is essentially bounded above by a maximum length $L_{max} \sim 2.16$ (full line). Moreover, the probability that the edge length exceeds the value 2 (dotted line) is about 10^{-3} .




Figure 3: **Time-frequency filtering — Synthetic signal example (1/2)**. Top left: spectrogram of a Hermite function embedded in white Gaussian noise ($\text{SNR} = 10$ dB). Bottom left: Delaunay triangulation constructed on the zeros of the spectrogram, with outlier edges (see text) highlighted in red. Bottom right: time-frequency domains obtained by concatenating adjacent Delaunay triangles with outlier edges, each domain being labeled by a color. Top right: masked spectrogram when retaining as domain the blue ribbon.




Figure 4: **Time-frequency filtering — Synthetic signal example (2/2)**. Top diagram: noisy observation. Bottom diagram: reconstructed waveform obtained by inverting the masked STFT of Figure 3 (top right), together with the noise-free Hermite function for a sake of comparison.




Figure 5: **Time-frequency filtering — Real data example (1/3)**. Left diagram: spectrogram of the benchmark “bat signal”. Right diagram: reassigned spectrogram.




Figure 6: **Time-frequency filtering — Real data example (2/3)**. Left diagram: Delaunay triangulation constructed on the zeros of the spectrogram, with outlier edges (see text) highlighted in red. Right diagram: time-frequency domains obtained by concatenating adjacent Delaunay triangles with outlier edges, each domain being labeled by a color.

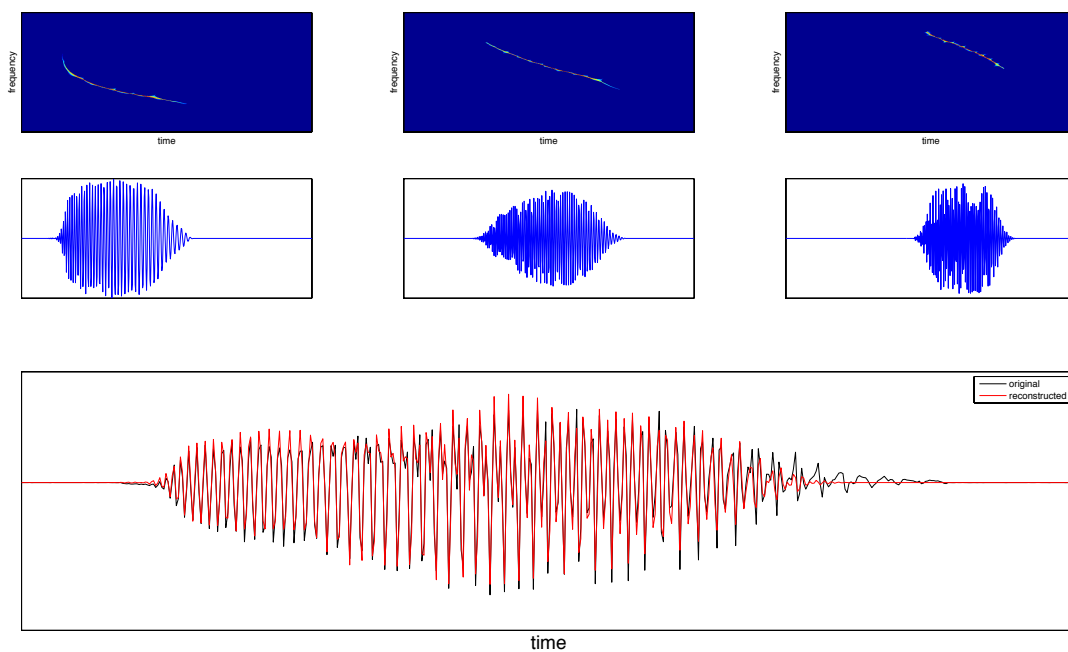


Figure 7: **Time-frequency filtering — Real data example (3/3)**. Top row: masked reassigned spectrograms of the 3 main components. Middle row: the corresponding waveforms obtained by inverting the respective masked STFTs. Bottom row: superimposition of the above 3 components, together with the original signal for a sake of comparison.