

Andreas Tzanis

*Department of Geophysics and Geothermy,
National and Kapodistrian University of Athens,
Panepistimiopoli,
Zografou 15784, Greece.
E-mail: atzanis@geol.uoa.gr*

**THE CURVELET TRANSFORM IN THE ANALYSIS OF 2-D
GPR DATA: SIGNAL ENHANCEMENT AND EXTRACTION OF
ORIENTATION-AND-SCALE-DEPENDENT INFORMATION**

Journal of Applied Geophysics, vol. 115, pp.145–170, 2015

DOI: <http://dx.doi.org/10.1016/j.jappgeo.2015.02.015>

Abstract

The Ground Probing Radar (GPR) is a valuable tool for near surface geological, geotechnical, engineering, environmental, archaeological and other work. GPR images of the subsurface frequently contain geometric information (constant or variable-dip reflections) from various structures such as bedding, cracks, fractures etc. Such features are frequently the target of the survey; however, they are usually not good reflectors and they are highly localized in time and in space. Their scale is therefore a factor significantly affecting their detectability. At the same time, the GPR method is very sensitive to broadband noise from buried small objects, electromagnetic anthropogenic activity and systemic factors, which frequently blurs the reflections from such targets. The purpose of this paper is to investigate the Curvelet Transform (CT) as a means of S/N enhancement and information retrieval from 2-D GPR sections, with particular emphasis on the recovery of features associated with specific temporal or spatial scales and geometry (orientation/dip).

The CT is a multiscale and multidirectional expansion that formulates an optimally sparse representation of bivariate functions with singularities on twice-differentiable (C^2 -continuous) curves (e.g. edges) and allows for the optimal, whole or partial reconstruction of such objects. The CT can be viewed as a higher dimensional extension of the wavelet transform: whereas discrete wavelets are isotropic and provide sparse representations of functions with point singularities, curvelets are highly anisotropic and provide sparse representations of functions with singularities on curves. A GPR section essentially comprises a spatio-temporal sampling of the transient wavefield which contains different arrivals that correspond to different interactions with wave scatterers in the subsurface (wavefronts). These are generally longitudinally piecewise smooth and transversely oscillatory, i.e. they comprise edges. Curvelets can detect wavefronts at different angles and scales because curvelets of a given angle and scale locally correlate with aligned wavefronts of the same scale.

The utility of the CT in processing noisy GPR data is investigated with software based on the Fast Discrete CT and adapted for use with a set of interactive driver functions that compute and display the curvelet decomposition and then allow the manipulation of data (wavefront) components at different scales and angles via the corresponding manipulation (cancellation or restoration) of their associated curvelets. The method is demonstrated with data from archaeometric, geotechnical and hydrogeological surveys, contaminated by high levels of noise, or featuring straight and curved reflections in complex propagation media, or both. It is shown that the CT is very effective in enhancing the S/N ratio by isolating and cancelling directional noise wavefronts of any scale and angle of emergence, sometimes with surgical precision and with particular reference to clutter. It can as successfully be used to retrieve waveforms of specific scale and geometry for further scrutiny, also with surgical precision, as for instance distinguish signals from small and large aperture fractures and faults, different phases of fracturing and faulting, bedding etc. Moreover, it can be useful in investigating the characteristics of signal propagation (hence material properties), albeit indirectly. This is possible because signal attenuation and temporal localization are closely associated, so that scale and spatio-temporal localization are also closely related. Thus, interfaces embedded in low attenuation domains will tend to produce sharp reflections and fine-scale localization. Conversely, interfaces in high attenuation domains will tend to produce dull reflections with broad localization.

Introduction

The Ground Probing Radar (GPR) has become a valuable, almost indispensable means of exploring near surface structures for geological, geotechnical, engineering, environmental, archaeological and other work. The purpose of this paper is to investigate methods of signal enhancement and geometrical information retrieval from GPR data, with particular emphasis placed on the problem of recovering features associated with specific temporal or spatial scales and geometry (orientation/dip).

GPR data essentially comprise recordings of the amplitudes of transient waves generated and detected by source and receiver antennae, with each source/receiver pair generating a data trace that is a function of time. An ensemble of traces collected sequentially along a scan line, i.e. a GPR section or B-scan, provides a spatio-temporal sampling of the wavefield which contains different arrivals that correspond to different interactions with wave scatterers (inhomogeneities) in the subsurface. All these arrivals represent wavefronts that are relatively smooth in their longitudinal direction and oscillatory in their transverse direction. Thus, GPR images frequently contain geometric (orientation/dip-dependent) information from point scatterers (e.g. diffraction hyperbolae), dipping reflectors (geological bedding, structural interfaces, cracks, fractures and joints), as well as other conceivable structural configurations. In geological, geotechnical and engineering applications, one of the most significant objectives is the detection of fractures, inclined interfaces and empty or filled cavities frequently associated with jointing/faulting; these types of target, especially fractures, are usually not good reflectors and are spatially localized.

At the same time, GPR data is notoriously susceptible to noise. For example, boulders, animal burrows, trees, tree roots, and other small scale objects and structures can cause unwanted reflections or scattering. Anthropogenic noise is worse and can include reflections from nearby vehicles, buildings, fences and power lines. Transmissions from cellular telephones, two-way radios, television, and radio and microwave devices may also cause noise. These types of reflections are only partially countered with shielded antennae while the interference by extraneous or reflected airwaves, critically refracted airwaves and ground waves cannot be easily suppressed during acquisition. Finally, there's systemic noise, frequently manifested in the form of ringing (antenna self-clutter). In many cases, the noise has definite directional characteristics (e.g. high-angle crossing clutter). Because the GPR source wavelet is tuned at a single operating frequency, the information returned by the subsurface structure is usually limited to a relatively narrow band around it (plus a tail due to dispersion) and the rest of the spectrum is swamped in noise. Raw GPR data frequently require post-acquisition processing, as they usually provide only approximate target shapes and distances (depths).

There are several methods to de-noise two dimensional data, focus on single or multiple scales and extract geometrical information. Almost all of them have been described in the excellent and comprehensive review of Jacques et al. (2011). They can be broadly classified into two categories: Directional Filters and Multi-Resolution Analysis (MRA).

Directional Filters (Directional Wavelets) are useful in many image processing tasks such as texture analysis, edge detection, image data compression, motion analysis, and image (signal) enhancement. Steerable Wavelets are an early form of Directional Filters (Freeman and Adelson, 1991; Simoncelli et al. 1992) and are closely related to the Gabor wavelets (e.g. Feichtinger and Strohmer, 1998, 2003; Lee, 2008). The Steerable Wavelets are based on directional second derivatives of a Gaussian function, while the Gabor wavelets are produced by a Gabor kernel, which is a product of an elliptical Gaussian and a complex plane wave. Tzanis (2013) discussed another class of two-dimensional directional filters, the B-

Spline Wavelet (BSW) Filters; these are built by sidewise arranging a number of one-dimensional B-spline wavelets to form a matrix, tapering the transverse direction with an orthogonal window function and rotating the resulting matrix to the desired orientation. The BSW and Gabor Filters were applied to the analysis of heavily contaminated GPR data and have been shown to be very effective. These methods may successfully process information at arbitrarily fine scales and single orientations but do not allow for a different number of directions at each scale. In order to obtain multidirectional representation of the data at each scale, it is necessary to apply the same filter rotated to different angles under adaptive control and combine the outputs (e.g. Freeman and Adelson, 1991, for Steerable Wavelets; Grigorescu et al. 2003, for Gabor Filters; Tzanis, 2013, for Gabor and B-Spline Wavelet Filters). With reference to GPR data, the results are usually remarkable but from a computational point of view, not particularly inexpensive.

Wavelet-based processing and analysis methods have been extensively applied to images (extraction of information, compression and de-noising). Two-dimensional geophysical data is very similar to an image if each data point is taken to be a pixel; in consequence, the same wavelet-based methods have been quite extensively adapted to the processing of seismic and potential-field geophysical data but only rarely to GPR data. One such effective wavelet-based method commonly used for the treatment of geophysical data, is Multi-Resolution Analysis (MRA, e.g. Chui, 1992; Mallat, 1999). This is the design method of most of the practically relevant discrete wavelet transforms and the justification of the fast wavelet transform. MRA allows a space $L^2(\mathbb{R}^2)$, i.e. an image, to be decomposed into a sequence of nested subspaces (images) $L^2(\mathbb{R}^2) \supset \dots \supset V_n \supset \dots \supset V_0 \supset \dots \{0\}$, arranged in order of increasing detail (scale), that satisfies certain self-similarity relations in time/space and scale/frequency, as well as completeness and regularity relations. This way, the MRA provides a means to suppress events of specific scales locally but leave the rest of the data unaffected in general. The MRA has been introduced to the processing of reflection seismics fairly recently, with most of the related studies attempting to develop efficient noise suppression procedures in a time-frequency sense. Inasmuch as the GPR method is effectively equivalent to reflection seismics, GPR data can be processed with similar techniques. The pertinent literature (seismic and GPR) is not rich, but is steadily growing in numbers and applications (e.g. Deighan and Watts, 1997; Leblanc et al. 1998; Miao and Cheadle, 1998; Matos and Osorio, 2002; Nuzzo and Quarta, 2004; Jeng et al. 2009).

MRA, as applied in the literature quoted above, is not suitable for processing orientation-dependent information. In a manner analogous to the realization that Fourier methods are not suitable (or adaptable) for all signal processing problems, which has consequently led to the introduction of the wavelet transform, it has also become apparent that wavelets could successfully deal only with phenomena that are generally isotropic except for local irregularities (i.e. associated with isolated singularities at exceptional points); wavelets are less than ideal in dealing with phenomena occurring on curves or sheets (i.e. with singularities on curves), as for instance, edges in a two-dimensional image, or reflections and wavefronts in a seismic or GPR record.

This problem has been addressed by advanced MRA-like algorithms that have appeared recently and are collectively referred to as the “*X-let Transform*”. These include *ridgelets* (Candès, 1999; Candès and Donoho, 1999), *wedgelets* (Donoho, 1999), *beamlets* (Donoho and Huo, 2002), *bandlets* (Mallat and Peyré (2007), *contourlets* (Do and Vetterli, 2005), *wave atoms* (Demanet and Ying, 2007), *surfacelets* (Lu and Do, 2007) and others. All these vary considerably in scope, properties and efficiency; details cannot be discussed herein but a comprehensive review can be found in Jacques et al. (2011). Other very effective and versatile approaches include the second generation *Curvelet Transform* (Candès and Donoho, 2003; 2003b; 2004a) and its derivatives/ extensions, the *Shearlet Transform* (e.g. Guo and Labate 2007) and the *Riplet Transform* (e.g. Xu et al. 2010). The latter methods are specifically designed

to associate scale with orientation; they yield optimally sparse representations of the data and have optimal reconstruction properties (see below). These very desirable characteristics of the Curvelet Transform lineage have motivated the present investigation, which, however, will focus only on the fundamental design: The Curvelet Transform (CT).

The CT is a multiscale and multidirectional expansion that formulates a sparse representation of the input data (a representation is sparse when it describes the data with a superposition of a small number of components). The roots of the CT are traced to the field of Harmonic Analysis, where curvelets were introduced as expansions for asymptotic solutions of wave equations (Smith, 1998; Candès, 1999). In consequence, curvelets can be viewed as primitive and prototype waveforms – they are local in both time/space and frequency/wavenumber and correspond to a partitioning of the 2D Fourier plane by highly anisotropic elements that obey the parabolic scaling principle, that is their width is proportional to the square of their length. Owing to their anisotropic shape, curvelets are well adapted to detect wavefronts at different angles and scales because curvelets at a given scale locally correlate with aligned wavefronts of the same scale (given that wavefronts are actually edges in an otherwise smooth image).

Based on such properties of the curvelets, the CT provides an optimally sparse representation of objects with edges – specifically of objects which are smooth except for discontinuities along general curves with bounded curvature (twice differentiable or C^2 -continuous curve). If f_m is the m -order curvelet approximation to an object $f(x_1, x_2) \in L^2(\mathbb{R}^2)$, Candès and Donoho (2004a) show that if the object is singular along a smooth C^2 curve, the approximation error is $\|f - f_m\|_2^2 = O(m^{-2}[\log m]^3)$ and is optimal in the sense that there is no other representation of the same order m , that can yield a smaller asymptotic error. This is far better than the $O(m^{-1})$ error afforded by the wavelet approximation. Optimal sparsity also implies that one can recover curved objects from noisy data by curvelet shrinkage (analogous to wavelet shrinkage) and obtain a mean squared error that is far better than what was affordable with more traditional methods (Candès et al. 2006). In addition to the “static” view of optimal sparsity, curvelets may also formulate a “dynamic” optimally sparse representation of wave propagators (e.g. Candès and Demanet, 2005; Sun et al. 2010): curvelets can model the geometry of wave propagation by translating the centre of the curvelet along the wave path. As Candès et al. (2006) point out, a physical interpretation of this result is that curvelets may be viewed as the representation of local plane waves, with sufficient frequency localization to behave like waves and sufficient spatial localization to simultaneously behave like particles. The “dynamic” approach to curvelets will not be considered herein.

Another cardinal property of the CT is optimal image reconstruction in severely ill-posed problems; The CT possesses microlocal properties, (i.e. localization properties dependent on both position and orientation), which render them particularly suitable for reconstruction problems with missing data. In general terms, depending on the acquisition geometry and noise characteristics, the curvelet expansion can be separated into a subset that can be recovered accurately, and one that cannot. The microlocal properties of the curvelets allow the former (recoverable) part to be reconstructed with accuracy similar to the accuracy that would be feasible if the data was complete and noise-free. To this effect, Candès and Donoho (2004b) show that for some statistical models which allow for C^2 objects to be recovered, there are simple algorithms based on the shrinkage of curvelet biorthogonal decompositions, which achieve optimal statistical rates of convergence, i.e. there is no other estimating procedure which can return a fundamentally better mean square error.

The main body of this paper is organized as follows: The 2nd Generation CT will first be introduced in Section 2. Inasmuch as the application of the method is still uncommon in Exploration Geophysics at large and GPR practice in particular, the introduction will be relatively rigorous; it will spare the reader from details and proofs but will elaborate on how curvelets are constructed and computed and how they

operate on data with wavefronts. An introduction to the practical implementation of filtering via the CT will be given in Section 3. This will be followed by a rigorous presentation of five applications (Section 4): the first three will concentrate on signal enhancement and information retrieval from data contaminated by high levels of noise; the last two will concentrate on information retrieval from data featuring straight and curved reflections in complex propagation media. Finally, a discussion of the results and comparisons with other methods developed with the same objective will conclude the presentation.

1 The Second Generation Curvelet Transform

1.1 The Continuous Curvelet Transform (CCT)

A curvelet frame is a wave packet frame on $L^2(\mathbb{R}^2)$ based on a second dyadic decomposition which, in effect, comprises an extension of the isotropic MRA concept to include anisotropic scaling and angular dependence (directionality) while maintaining rotational invariance. In order to construct the curvelet frame, a template (basic) curvelet is required, which will generate it by translation, dilation *and* rotation. The elements of the curvelet family will provide a partition (tiling) of the two-dimensional Fourier plane.

Consider a function $f(\mathbf{x}) \in L^2(\mathbb{R}^2)$ with $\mathbf{x} = [x_1 \ x_2]^T$ and x_j represent space or/and time. Next, consider the Fourier transform pair $f(\mathbf{x}) \leftrightarrow F(\boldsymbol{\xi})$ where $\boldsymbol{\xi} = [\xi_1 \ \xi_2]^T$ and ξ_j represent frequency or/and wavenumber. Let $r=|\boldsymbol{\xi}|$ be the radial coordinate and $\theta = \arctan(\xi_1, \xi_2)$ be the azimuthal coordinate in the ξ -plane. Finally, consider a partition of the polar coordinate plane in concentric annuli (*coronae*) according to $2^{j-1} \leq r \leq 2^{j+1}$, with each corona further partitioned into angular sectors according to $\angle(\theta) \leq 2^{-j/2}$, so as to generate *polar wedges* (Fig. 1a). The number of the wedges N_j increases like $\sqrt[3]{(1/\text{scale})}$ and doubles in every second ring so that the width of the wedges is proportional to the square of their length: this is the so called *parabolic scaling* (Fig. 1a).

Now consider a *radial* window $W(r)$ and an *angular* window $V(t)$, which must be smooth, nonnegative and real-valued. The support of W is $r \in (1/2, 2)$ and it must obey the *admissibility condition* $\sum_{j=-\infty}^{\infty} |W(2^{-j}r)|^2 = 1$ for $r > 0$; the support of V is $t \in (-2\pi, 2\pi]$ and the corresponding admissibility condition $\sum_{l=-\infty}^{\infty} V^2(t - 2\pi l) = 1$ for $t \in \mathbb{R}$. There is no other theoretical limitation on the nature of the windows, which means that they can be wavelets (e.g. as in Ma and Plonka, 2010).

The dilated basic curvelet in polar coordinates is

$$\Phi_{j,l,k=0}(r, \theta) = 2^{-3j/4} W(2^{-j}r) V\left(\frac{2^{\lfloor j/2 \rfloor} \theta}{2\pi}\right), \quad r \geq 0, \quad \theta \in [0, 2\pi), \quad j \in \mathbb{N}_0, \quad (1)$$

with $\lfloor \cdot \rfloor$ denoting the floor operator (integer part). Evidently, $W(2^{-j}r)$ isolates ξ -values in the corona $(2^{j-1}, 2^{j+1})$. Likewise, $V(2^{\lfloor j/2 \rfloor} \theta)$ isolates ξ -values in the angular sector $(-\pi 2^{-\lfloor j/2 \rfloor}, \pi 2^{-\lfloor j/2 \rfloor})$. For increasing j and decreasing scale $2^j \in (0, 1]$, the breadth of W is growing and the width of V is shrinking, so that the wedges $\Phi_{j,0,0}(r, \theta)$ become longer (see Fig. 1a and 1c): this is the effect of parabolic scaling.

The complete curvelet family in the ξ -domain is generated from $\Phi_{j,0,0}$ by dilation, rotation and translation according to:

$$\Phi_{j,l,k}(r, \theta) = \Phi_{j,0,0}\left(\mathbf{R}_{\theta_j} \boldsymbol{\xi}\right) \cdot e^{-i\langle \mathbf{x}_k^j, \boldsymbol{\xi} \rangle}, \quad (2a)$$

or, in terms of the radial and angular windows,

$$\Phi_{j,l,k}(r, \theta) = 2^{-3j/4} W(2^{-j} r) \cdot V\left(\frac{2^{\lfloor j/2 \rfloor}}{2\pi}(\theta - \theta_l)\right) \cdot e^{-i\langle \mathbf{x}_k^j, \xi \rangle}, \quad (2b)$$

where, $\theta_{j,l} = 2\pi l \cdot 2^{-\lfloor j/2 \rfloor}$, $l = 0, 1, 2, \dots$, $0 \leq \theta_{j,l} < 2\pi$, is a sequence of equi-spaced rotation angles whose number varies proportionally to $1/\sqrt{\text{scale}}$,

$$\mathbf{R}_{\theta_{j,l}} = \begin{bmatrix} \cos \theta_{j,l} & \sin \theta_{j,l} \\ -\sin \theta_{j,l} & \cos \theta_{j,l} \end{bmatrix}, \quad \mathbf{R}_g^{-1} = \mathbf{R}_g^T = \mathbf{R}_{-\theta},$$

is the rotation by $\theta_{j,l}$ radians and

$$\mathbf{x}_k^{j,l} = \mathbf{R}_{\theta_{j,l}}^{-1} \begin{bmatrix} 2^{-j} & 0 \\ 0 & 2^{-j/2} \end{bmatrix} \begin{bmatrix} k_1 \\ k_2 \end{bmatrix}$$

are scaled positions with $k_1, k_2 \in \mathbb{Z}^2$ representing the translation parameters. As evident in Eq. (2), the angular window V now isolates ξ -values in the angular sector $-\pi 2^{-\lfloor j/2 \rfloor} \leq \theta - \theta_{j,l} \leq \pi 2^{-\lfloor j/2 \rfloor}$. Moreover, the support of $\Phi_{j,l,k}$ does not depend on the translation parameters at all!

In the x -domain, the curvelet comprises a waveform defined by the Fourier Transform of Φ :

$$\varphi_{j,l,k}(\mathbf{x}) = \varphi_{j,0,0}(\mathbf{R}_{\theta_{j,l}}(\mathbf{x} - \mathbf{x}_k^{j,l})). \quad (3)$$

Inasmuch as the support of $\Phi_{j,l,k}$ is independent of \mathbf{k} , the frequency localization of $\Phi_{j,l,0}$ is such that $\varphi_{j,l,0}(\mathbf{x})$ will decay rapidly outside a rectangle of size $2^{-j} \times 2^{-j/2}$ with its major axis perpendicular to the polar angle $\theta_{j,l}$ (e.g. Fig. 1b and 1d). Thus, the effective size of $\varphi_{j,l,k}(\mathbf{x})$ will also obey the parabolic scaling relationships: $\text{length} \approx 2^{-j/2}$, $\text{width} \approx 2^{-j} \Rightarrow \text{width} \approx \text{length}^2$. Moreover, $\Phi_{j,0,0}$ is by construction supported away from the ξ_2 axis (where $\xi_1 = 0$) but near the ξ_1 axis (where $\xi_2 = 0$). Therefore, $\varphi_{j,0,0}(\mathbf{x})$ will oscillate in the x_2 -direction and will be low-pass in the x_1 -direction. As evident from Eq. (2), these oscillation properties are preserved during rotation and translation (e.g. Fig. 1d). Then, at any scale 2^j , a curvelet will be enveloped by a ridge with effective length $2^{-j/2}$ and width 2^{-j} and will oscillate in a direction perpendicular to that ridge (see Fig. 1d).

A last matter to be considered is the ‘‘hole’’ arising in the ξ -plane around zero, since the rotations of the basic curvelets are defined only for the scales 2^j . For a complete covering of the ξ -plane one must define a low-pass element W_{j_0} which is supported on the unit circle and is non-directional (isotropic). This (coarsest scale) element obeys

$$|W_{j_0}(r)|^2 + \sum_{j>j_0} |W(2^{-j} r)|^2 = 1, \quad (4)$$

so that coarsest scale curvelet will be

$$\Phi_{j_0}(\xi) = 2^{-j_0} W_{j_0}(2^{-j_0} |\xi|),$$

and will admit a x -domain representation of the form

$$\varphi_{j_0,k}(\mathbf{x}) = \varphi_{j_0}(\mathbf{x} - 2^{-j_0} \mathbf{k}).$$

Based on the above construction principles, a *curvelet coefficient* comprises the inner product between $f(x_1, x_2)$ and a curvelet $\varphi_{j,l,k}, j > j_0$:

$$c(j, l, k) = \langle f, \varphi_{j,l,k} \rangle = \int_{\mathbb{R}^2} f(\mathbf{x}) \overline{\varphi_{j,l,k}(\mathbf{x})} d\mathbf{x}. \quad (5)$$

The inverse continuous curvelet transform will then be defined by the reconstruction rule

$$f(\mathbf{x}) = \sum_{j \geq j_0, l, k} c(j, l, k) \varphi_{j,l,k}(\mathbf{x}), \quad (6)$$

while the Parseval condition $\sum_{j,l,k} |c(j,l,k)|^2 = \|f\|_{L^2(\mathbb{R}^2)}^2$ will hold for all $f \in L^2(\mathbb{R}^2)$ ensuring a *tight frame* property. As a result of the parabolic scaling, the curvelet frame is an optimally sparse representation for functions with singularities along curves but otherwise smooth (Candès and Demanet, 2005).

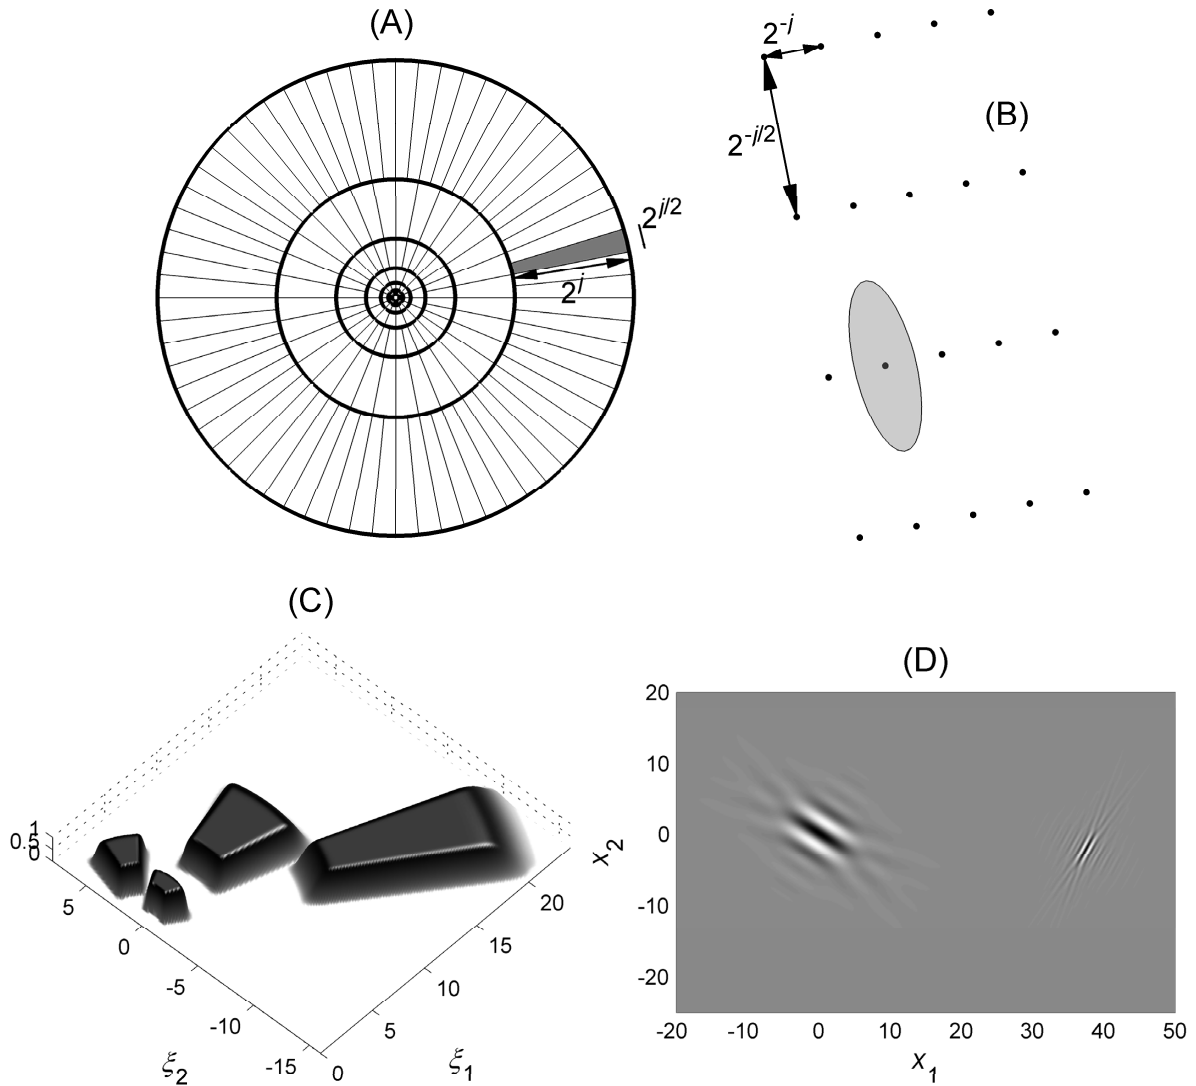


Figure 1. (a) Tiling of the ξ -plane in polar coordinates with parabolic scaling. The shaded area represents a wedge supporting a curvelet. (b) Schematic representation of a Cartesian grid in the x -domain, associated with a ξ -domain wedge like the shaded one shown in Fig. 1a. Due to the duality between the two domains, the spacing and scaling of the x -domain curvelet (represented by the ellipse) is also parabolic. (c) Some ξ -domain curvelets in perspective view: from left to right they are $\{j=1, l=0\}$, $\{j=2, l=2\}$, $\{j=3, l=0\}$ and $\{j=4, l=22\}$. (d) The x -domain curvelets $\{j=2, l=2\}$ (left) and an arbitrarily translated version of $\{j=4, l=22\}$ (right).

1.2 The Discrete Curvelet Transform (DCT)

The formulation of the CCT is not particularly suitable for measured 2-D data which is usually obtained in the form of rectangular (Cartesian) arrays, because circular coronae and rotation are not easily adaptable to Cartesian geometries. In response, the inventors of the curvelet transform also developed discrete formulations which are suitable for Cartesian data arrays while being faithful to the mathematical formalism of the CCT.

In these formulations, the circular coronae are replaced by rectangular “coronae” (concentric rectangular annuli) and the rotations are replaced by shearing, so as to generate a *pseudo-polar* grid like the one shown in Fig. 2a. Thus, the Cartesian equivalent of the radial window W is

$$\tilde{W}_j(\xi) = \sqrt{U_{j+1}^2(\xi) - U_j^2(\xi)}, \quad j \geq 0,$$

where $U(\xi) = u(2^{-j}\xi_1) \cdot u(2^{-j}\xi_2)$ and u is a low-pass one-dimensional window that obeys $0 \leq u \leq 1$ and vanishes for all $\xi \notin [-2, 2]$. Under this definition and on including the isotropic window $\tilde{W}_{j_0}(\xi)$, of the coarsest scale, the Cartesian analogue of the admissibility condition (4) reads

$$\tilde{W}_{j_0}^2(\xi) + \sum_{j>j_0} \tilde{W}_j^2(\xi) = 1 \quad (7)$$

and can be shown to hold for all ξ . The angular window V is now taken to be

$$V_j(\xi) = V \left(\begin{array}{c} 2^{\lfloor j/2 \rfloor} \xi_2 \\ \xi_1 \end{array} \right).$$

Following the above definitions, the basic curvelet in Cartesian coordinates is

$$\tilde{\Phi}_{j,l=0,k=0}(\xi) = 2^{-3j/4} \cdot \tilde{W}_j(\xi) \cdot V_j(\xi)$$

and is supported on a *trapezoidal* wedges bounded as $\{(\xi_1, \xi_2) : 2^j \leq \xi_1 \leq 2^{j+1}, -2^{-j/2} \leq \frac{\xi_2}{\xi_1} \leq 2^{-j/2}\}$, isolating the ξ -values included therein.

In order to map the basic Cartesian curvelet onto different orientations, instead of equi-spaced angles define a set of *equi-spaced slopes* $\tan \theta_l = l \cdot 2^{-\lfloor j/2 \rfloor}$, $l = -2^{\lfloor j/2 \rfloor}, \dots, 2^{\lfloor j/2 \rfloor} - 1$ and define

$$\tilde{\Phi}_{j,l,k=0}(\xi) = \tilde{\Phi}_{j,0,0}(\mathbf{S}_{\theta_l} \xi) \quad (8a)$$

where

$$\mathbf{S}_{\theta_l} = \begin{bmatrix} 1 & 0 \\ -\tan \theta_l & 1 \end{bmatrix}$$

is the *shear* matrix and where the admissibility condition is upheld. Eq. (8a) can be written in terms of the windows \tilde{W} and V as

$$\tilde{\Phi}_{j,l,k=0}(\xi) = 2^{-3j/4} \cdot \tilde{W}_j(\xi) \cdot V_j(\mathbf{S}_{\theta_l} \xi) = 2^{-3j/4} \cdot \tilde{W}_j(\xi) \cdot V_j \left(\begin{array}{c} 2^{\lfloor j/2 \rfloor} \xi_2 - l \\ \xi_1 \end{array} \right) \quad (8b)$$

and clearly isolates ξ -values in the wedge $\{(\xi_1, \xi_2) : 2^j \leq \xi_1 \leq 2^{j+1}, -2^{-j/2} \leq \frac{\xi_2}{\xi_1} - \tan \theta_l \leq 2^{-j/2}\}$.

Constructions such as those implied by Eq. (8) are illustrated in Fig. 2b. In order to define the Cartesian analogues of the family $\Phi_{j,l,0} = \Phi_{j,0,0}(\mathbf{R}_{\theta_j} \xi)$ described in Section 2.1, it is sufficient to complete the curvelets $\tilde{\Phi}_{j,l,k=0}(\xi)$ by symmetry with respect to the origin and rotation by $\pm\pi/2$ radians. The complete family $\tilde{\Phi}_{j,l,k=0}(\xi)$ generates the concentric partitioning whose geometry is shown in Fig. 2a.




Figure 2. (A) Pseudo-polar partitioning (tiling) of the ξ -plane in Cartesian coordinates with trapezoidal wedges. The inner (coarsest-scale) isotropic partition corresponds to $j = 1$. The indexing of the wedges (l) counts clockwise from the top-left corner of each scale. The solid black partitions at $j=3, l=5$ and $j=3, l=15$ indicate the right-hand side trapezoidal wedges that supports the curvelets shown in (B). The gray partitions indicate the wedges supporting the curvelets shown in Fig. 3.

In a final step, on introducing the translation parameters

$$\tilde{\mathbf{x}}_k^{j,l} = \mathbf{S}_{\theta_l}^{-T} \begin{bmatrix} 2^{-j} & 0 \\ 0 & 2^{-j/2} \end{bmatrix} \begin{bmatrix} k_1 \\ k_2 \end{bmatrix} = \mathbf{S}_{\theta_l}^{-T} \mathbf{b}_k^j,$$

and using Eq. (8b), the Cartesian equivalent of Eq. (2) can be written thus:

$$\tilde{\Phi}_{j,l,k}(\xi) = \tilde{\Phi}_{j,0,0}(\mathbf{S}_{\theta_l} \xi) \cdot e^{-i\langle \tilde{\mathbf{x}}_k^{j,l}, \xi \rangle}.$$

As with its CCT counterpart, the x -domain representation of the Cartesian curvelet comprises a waveform defined by the Fourier Transform of $\tilde{\Phi}$ and is:

$$\tilde{\varphi}_{j,l,k}(\mathbf{x}) = \tilde{\varphi}_{j,0,0}(\mathbf{S}_{\theta_l}^T (\mathbf{x} - \tilde{\mathbf{x}}_k^{j,l})).$$

It follows that the discrete Cartesian counterpart of the curvelet coefficients is:

$$\tilde{c}(j, l, k) = \langle f, \tilde{\varphi}_{j,l,k} \rangle = \sum_{x_1, x_2} f(\mathbf{x}) \overline{\tilde{\varphi}_{j,l,k}(\mathbf{x})}. \quad (9)$$

Finally, the inverse discrete curvelet transform is defined by a reconstruction formula analogous to Eq. (6):

$$f(\mathbf{x}) = \sum_{j \geq j_0, l, k} \tilde{c}(j, l, k) \tilde{\varphi}_{j,l,k}(\mathbf{x}), \quad (10)$$

The implementation of Eq. (9) is not straightforward; the difficulty is that on using the two-dimensional Fourier Transform to obtain $\tilde{\varphi}_{j,l,k}(\mathbf{x})$ from $\tilde{\Phi}_{j,l,k}(\xi)$, one would need to evaluate the FFT on the sheared grid $\tilde{\mathbf{x}}_k^{j,l}$, where it cannot be applied. In response, two indirect solutions have been developed by Candés et al (2005): the *unequi-spaced FFT* method and the *wrapping* method.

In the USFFT method, the translation grid is rotated so as to align with the orientation of the curvelet. To do this, at each scale/angle pair (j, l) the implementation uses a non-standard interpolation scheme to

obtain sampled values of $F(\xi)$ over the support of $\tilde{\Phi}_{j,l,k}(\xi)$. The inverse transform also uses conjugate gradients iteration to invert the interpolation step. In consequence, the USFFT method has a rather higher computational cost than the wrapping algorithm.

In the wrapping method the translation grid remains untilted and the same for every orientation albeit each curvelet is assigned with its proper angle. The curvelet coefficients are taken as per Eq. (9), except that $\tilde{\mathbf{x}}_k^{j,l}$ is replaced by $\mathbf{b}_k^j = (k_1 2^{-j}, k_2 2^{-j/2})$ with \mathbf{b} taking values on a rectangular grid. In this case, however, it is apparent that the window $\tilde{\Phi}_{j,l,k}(\xi)$ cannot fit into a rectangle of size $2^j \times 2^{j/2}$ to which an inverse FFT could be applied. The wrapping algorithm addresses this problem by periodizing the windowed ξ -domain coefficients $\tilde{C}_{j,l}(\xi) = F(\xi)\tilde{\Phi}_{j,l,k}(\xi)$ and re-indexing $\tilde{C}_{j,l}$ by wrapping it around a rectangle centred at the origin and approximately equal to $2^j \times 2^{j/2}$. If $F(\xi)$ is unity, then $\tilde{C}_{j,l}(\xi) = \tilde{\Phi}_{j,l,k}(\xi)$ and a x -domain representation of the curvelet can be obtained by inverse FFT, otherwise the curvelet coefficients are obtained by inverse FFT of $\tilde{C}_{j,l}(\xi)$. Fig. 3a presents examples of curvelets computed with the wrapping method at different scales, orientations and (arbitrary) translations.

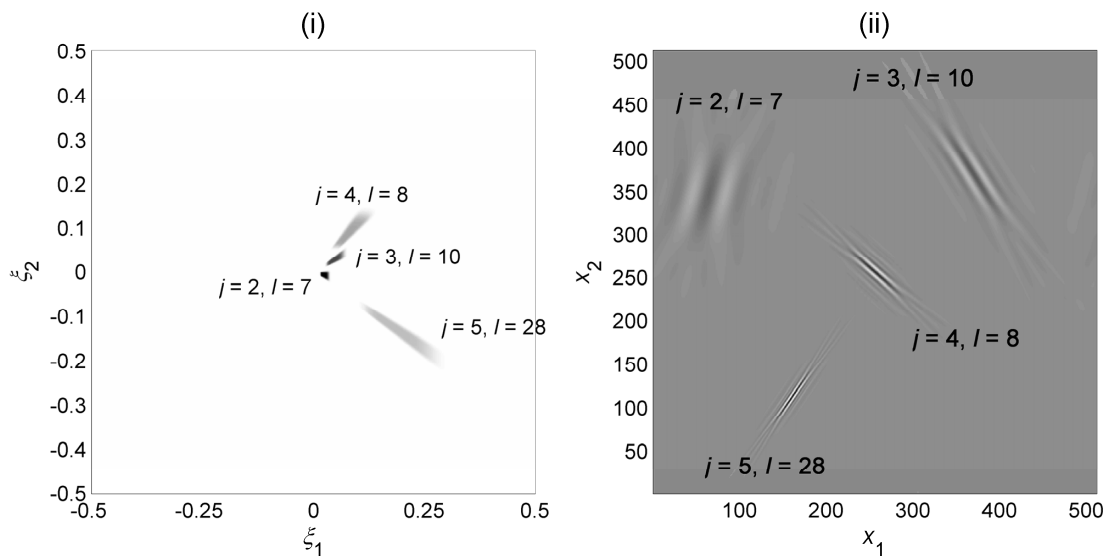


Figure 3a. (i) Amplitudes of complex ξ -domain curvelets at different scales and orientations. The trapezoidal wedges supporting these curvelets are shown in Fig. 2a. The indexing (l) increases clockwise from the top-left corner of each scale. (ii) Arbitrarily translated x -domain curvelets, corresponding to the ξ -domain curvelets of the left panel (i).

Fig. 3b shows how curvelets interact with data and elongate curved objects in particular. Specifically, panel (i) illustrates a contrived 512×512 data set featuring a wavy up-dipping set of intermittent reflections. The data was decomposed into six scales with the second coarsest scale ($j = 2$) comprising 24 angles (six per quadrant) and the number of angles doubling in every second scale. As evident from the previous discussion, curvelets are elongate and slender waveforms with length proportional to $2^{j/2}$ and width proportional to 2^j ; they oscillate in their transverse direction and are low-pass in their longitudinal direction. Curvelets can interact with curved objects in three ways:

1. When the curvelet and the object intersect while aligned parallel to their longitudinal directions, the transverse oscillatory part of the curvelet will locally match the same-scale component of the object and will encode the information in the corresponding set of curvelet coefficients which will have *significant* amplitudes.

2. When the curvelet and the object intersect at an arbitrary angle, the matching of the same-scale content will be imperfect and information will be lost to the curvelet's low-pass longitudinal action: the curvelet coefficients will have small amplitudes.
3. When the curvelet and the object do not intersect, the coefficients will be near zero.

Due to the unique orientational characteristics of this data set, only a handful of coefficients have noteworthy amplitudes. As an example, in Fig. 3b-ii the left panel illustrates the coefficients $\{j = 4, l = 5\}$ and the right panel a partial reconstruction based on these coefficients only! The associated curvelet has a slope of -110° and is almost perfectly aligned with the main trend of the "reflections" shown in Fig. 3b-i. In consequence, it extracts a clear and strong component of the signal. Conversely, the left panel of Fig. 3b-iii illustrates the coefficients $\{j = 4, l = 7\}$ and the right panel the corresponding partial reconstruction of the data. The associated curvelet has a slope of -94° and is intercepting the reflections at an angle of 14° ; evidently, it cannot match a signal component of any significant amplitude and the coefficients and reconstruction have very low amplitudes. Finally, the left panel of Fig. 3b-iv illustrates the coefficients $\{j = 3, l = 5\}$ and the right panel the corresponding partial reconstruction. The associated curvelet is perfectly aligned with the main reflections but has the wrong scale: it extracts only a very weak background component associated with the longer ξ -components of the data. It turns out that the data can be reconstructed to within 96% from the coefficients $\{j = 3, l = 4-6\}$. This demonstrates how curvelets may assist in the retrieval of scale-and-orientation dependent information from GPR data.

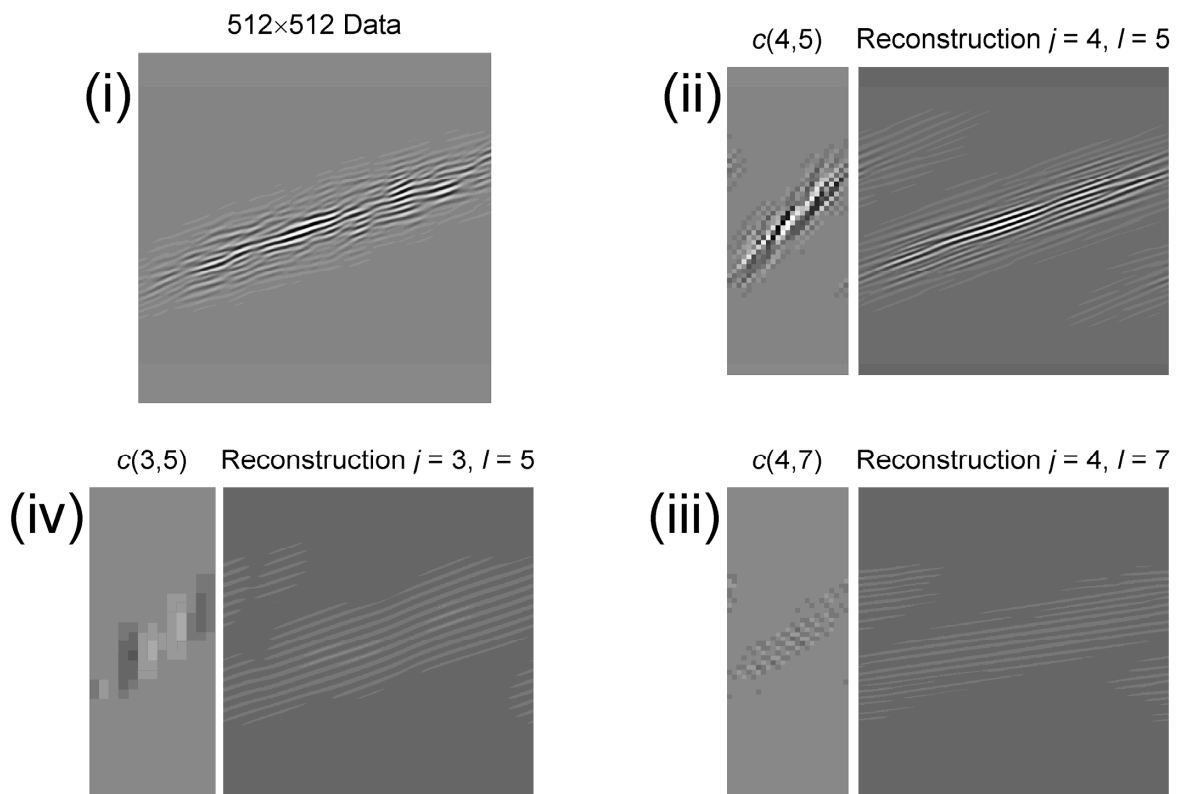


Figure 3b. A demonstration of data and curvelet interactions. **(i)** The data comprise a 512×512 matrix featuring only a set of wavy intermittent up-dipping reflections. **(ii)** The coefficients, (left) and a partial reconstruction of the data (right) generated by the curvelet $\{j = 4, l = 5\}$. **(iii)** As in (ii) but for $\{j = 4, l = 7\}$. **(iv)** As in (ii) and (iii) but for $\{j = 3, l = 5\}$.

An additional significant observation to be made in Fig. 3b is the following: The images of the curvelet coefficients and reconstructions, *all* contain a “side lobe” structure above the left and below the right end of the reconstructed up-dipping synthetic waveform. It must be emphasized that these are not artefacts but proper features of the reconstruction. Note that the synthetic waveform shown in Fig. 3b-i is actually truncated at the left and right ends of the image, i.e. it is incomplete. However, because the main part of the waveform contains sufficient information, the reconstruction rebuilds the missing ends and because there is no room to fit them in the 512×512 data matrix, it folds them around the edges. This is a powerful demonstration of the optimal reconstruction afforded by the Curvelet Transform and empirical justification of the statement made in the Introduction, that “... *the microlocal properties of the curvelets allow the recoverable part [of the data] to be reconstructed with accuracy similar to the accuracy that would be feasible if the data was complete and noise-free*”.

2 GPR data Analysis with the Curvelet Transform.

The DCT has been implemented in the software package CurveLab, written by E. Candés, L. Demanet and L. Ying and available at <http://www.curvelet.org>. The package contains Matlab and C++ implementations of both the USFFT and Wrapping methods. In the work presented herein, only the wrapping method has been implemented for being computationally more efficient. The utility of the CT will be demonstrated via a piece of software specifically developed to facilitate the interactive application of the DCT and to visualize and control the process; a fully functional version of this application is already available in the matGPR software (e.g. Tzanis, 2010), currently located at <http://users.uoa.gr/~atzanis/matgpr/matgpr.html>.

For a given dataset, the software prompts for the number of scales, with $j = 1$ corresponding to the (coarsest-scale) inner isotropic partition. It also inquires for the number of angles at the second coarser scale ($j = 2$), i.e. the first scale at which an angular decomposition is computed (the number of angles doubles in every second scale). The DCT is then computed with the option that yields real-valued curvelet coefficients: at any given scale, if L is the number of angles, then there will be $L/2$ ‘cosine’ coefficients stored with indices $1 \dots L/2$, as well as $L/2$ ‘sine’ coefficients stored by symmetry of $L/2$ with indices $L/2+1 \dots L$. The coefficients are indexed in a clockwise sense starting at the top-left corner of the north quadrant and are arranged so that the indices $l = 1 \dots L/2$ of the cosine coefficients span the north and east quadrants and the indices $l = L/2+1 \dots L$ of the sine coefficients span the south and west quadrants. Thus, each element of the data reconstruction rule specified by Eq. (10) can be completely specified by the pair $(l, L/2 + l)$ of sine and cosine coefficients.

The finest level scale extracts the highest frequency/ longest wavenumber content. The design of appropriate basis functions at the outermost dyadic corona is not straightforward due to difficulties relating to issues of over- or under- sampling at the high- ξ end of the spectrum (Candés et al. 2006). One apparent solution is to assign wavelets instead of curvelets to the finest scale and treat it with isotropic Cartesian windows $\tilde{W}_{j_e}(\xi)$ constructed so that $\tilde{W}_{j_e}^2(\xi)$, together with the curvelet windows, form a partition of unity as in Eq. (7). The wavelet coefficients can, then, be obtained by inverse FFT. This approach is simple but not consistent with the idea of directional basis elements at the finest scale. Proper curvelet-based remedies for tiling the finest scale have been proposed by Candés et al. (2006). Note, however, that GPR data is usually oversampled, at least with respect to the central frequency of the Tx antenna. In consequence, there is little (if any) useful information at the high frequency/ long

wavenumber scales which usually comprise random noise. In response, the wavelet-based remedy for the finest level scale was also adopted in this implementation of the DCT.

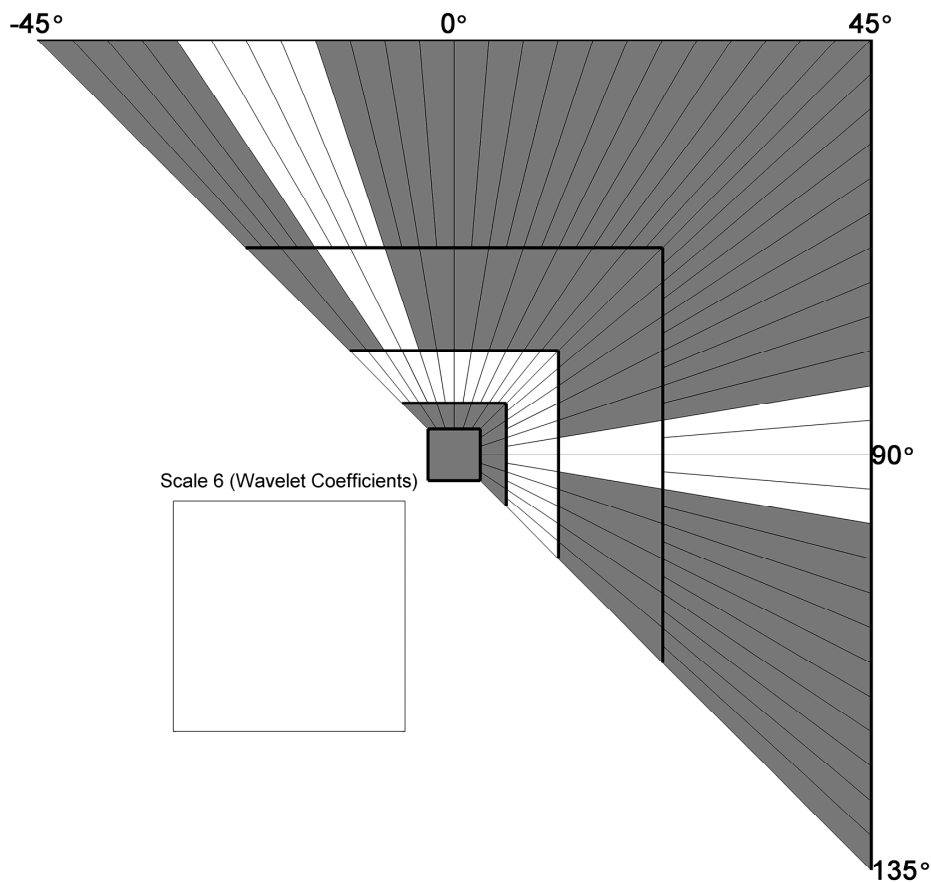


Figure 4. The upper-right diagonal (north and east quadrants) of the pseudo-polar tiling of the ξ -plane for a 512-by-512 data matrix decomposed into six scales and 24 angles at the second coarser scale ($j=2$). The tiling is drawn in data matrix coordinates and all angles and slopes refer to matrix coordinates (see text for details). Each trapezoidal wedge, as well as the rectangles corresponding to $j=1$ (central) and $j=6$ (lower left), is associated with a corresponding set of sine and cosine coefficients and functions as a “graphical switch” whose state can be toggled by pointing and clicking. The “On” state is indicated with shading and the corresponding coefficients are included in a reconstruction of the input data. The “Off” state is blank and the corresponding coefficients are excluded from a reconstruction.

After the DCT is computed, the upper-right diagonal of the (pseudo-polar) tiling of the ξ -plane is displayed as shown in Fig. 4 for a hypothetical 512-by-512 data matrix decomposed into six scales and 24 angles at the second coarser scale. The lower-left diagonal does not contain independent information and is redundant (see above). It is *very* important to clarify that the tiling is drawn in data matrix coordinates and not in physical coordinates of wavenumber and frequency. Likewise, all the angles (slopes) indicated in the drawing refer to data matrix coordinates. To see why, consider that in the discrete (digital) implementation of the CT, the curvelet decomposition can only be configured in matrix coordinates (just like in any other digital operation). Accordingly, in the ξ_2 -axis representing frequency, the Nyquist is $n_2/2$, $n_2 \in \mathbb{N}$, and the frequency range varies with integer steps in the interval $[-n_2/2, n_2/2]$. Likewise, in the ξ_1 -axis representing wavenumber, the Nyquist is $n_1/2$, $n_1 \in \mathbb{N}$, and the wavenumber varies with integer steps in $[-n_1/2, n_1/2]$. It follows that the dimensions of the data matrix are very important in determining the angular partitioning of each scale. With $n_2 \times n_1$ denoting the order of rows and columns, a

512×512 matrix has different angular partitioning than a 512×1024 or a 1024×512 matrix. In the first case the partitioning along the frequency and wavenumber axes is identical. In the second case the partitioning is more detailed along the wavenumber axis and in the third case it is more detailed along the frequency axis. A corollary of this point is that in some cases it may be advisable to appropriately resample the data so as to better focus on its temporal or spatial details.

Each trapezoidal wedge shown in Fig. 4 is a graphical object that represents the support of the curvelet used to generate its associated set of coefficients: at each scale the l^{th} graphical object (wedge) is assigned with the indices of the corresponding cosine (l) and sine ($L/2 + l$) coefficients, while it is also set up to function as a virtual switch whose On/Off state is controlled by the screen pointing device (e.g. the mouse). All the wedges are initially displayed in their ‘On’ state, which is indicated by shading. The ‘Off’ state is indicated by a blank wedge. The coarsest level isotropic partition is displayed as a rectangle at the centre of the tile. The finest level scale is displayed separately at the space that would have been occupied by the redundant lower-left diagonal and illustrates the associated set of wavelet coefficients. Both the coarsest and finest scale objects function as virtual switches just like the trapezoidal wedges.

It is possible to decide which coefficients to include (exclude) from a processed (reconstructed) version of the input data by pointing and clicking inside a trapezoidal wedge or rectangle. If the switch is ‘On’, this will *shrink* the associated sets of cosine and sine coefficients by a factor $\in [0, 1]$ and will toggle its state to ‘Off’, blanking it out as shown in Fig. 4. This is a *manual curvelet shrinkage* approach. The shrinking factor is specified by the analyst and can be changed during runtime. The default value is zero (complete negation of the selected coefficients) and all applications to be shown herein use the default value. The shrunk (or negated) coefficients may be restored by clicking inside a blank switch, in which case its state is reset to ‘On’ (shaded). It is also possible to toggle entire scales or angular subsets of a scale, by using appropriate GUI controls (not shown in Fig. 4). A partial or whole reconstruction of the data can be computed at any time, with GUI controls also not shown.

The application of the DCT to the analysis of GPR data will now be demonstrated with a complex synthetic radargram. Non-random noise is very difficult to simulate realistically and will not concern this example; the problem will be addressed in Section 4, in association with the treatment of real (observed) radargrams. Moreover, because a basic idea (and motivation) for implementing the CT is the manipulation and extraction of geometrical information, the analysis will be mainly focused on its angular resolution capability. The model used for generating the synthetic example is illustrated at the top of Fig. 5a and is summarized in Table 1. It consists of seven closely spaced geometrical objects with diverse material properties and orientations, embedded in a background with “average” material properties. The objects are so arranged, as to generate overlapping of wavefronts and clutter. This is useful in clarifying the reconstructive capacity, as well as the extent and limitations of the angular discriminative capacity affordable by the CT. It will also be useful in pointing out and discussing the output of the analysis with particular attention to the case of highly anisotropic filtering.

The zero-offset synthetic radargram was computed for a nominal central frequency of 400MHz with the time domain Finite Difference algorithm of Irving and Knight (2006) as implemented in the matGPR software (Tzanis, 2010). The synthetic section is shown in the middle panel of Fig. 5a after deconvolving the surface reflection; it comprises a 256 sample × 512 trace matrix with a sampling rate of 0.2149 ns and a trace spacing of 0.007436 m. The principal wavefronts (reflections) corresponding to the objects of Fig 5a are clearly discernible and pointed to. In addition, minor irregularities in the shape of objects 6 and 7, reverberation between cracks and scattering at the tips and intersections of the cracks also generates a significant amount of lower amplitude secondary wavefronts that take the form of crossing clutter.

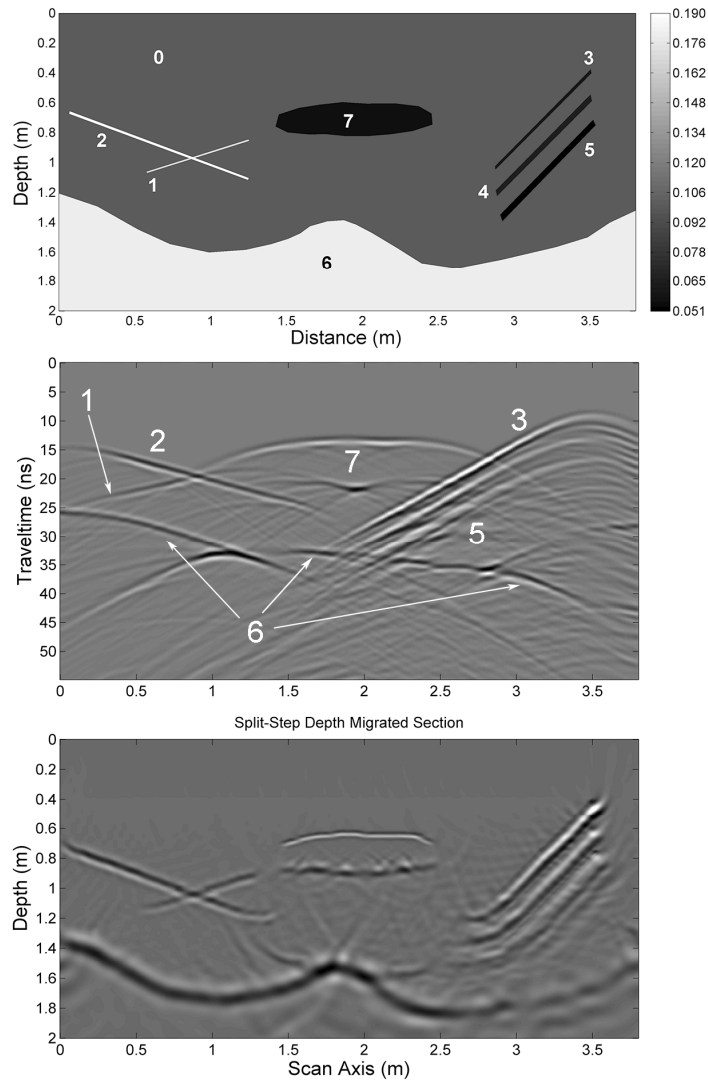


Figure 5a. Top: Synthetic structural model used for demonstrating the application of the Curvelet Transform. The objects comprising the model are shaded according to the non-dispersive term of the phase velocity and are described in Table 1. **Middle:** Synthetic zero-offset radargram generated from the model shown at the top panel with the FDTD solver of Irving and Knight (2006); the central frequency is 400MHz. **Bottom:** The synthetic radargram after 2-D depth migration with the split-step method (see text for details).

Table 1. Description and properties of the objects comprising the synthetic model of Fig. 5a. The velocity and quality factors are calculated after Bano (1996)

Object	Description	Dip (°)	σ (S/m)	K	μ_r	Velocity (m/ns)	Q
0	Background	–	10^{-3}	9	1	0.1	100
1	Short, thin dielectric crack.	-17°	10^{-3}	3	1	0.17	34
2	Dielectric crack.	21°	10^{-3}	2.5	1	0.19	28
3	Thin crack with moist argillaceous material	-45°	5×10^{-3}	25	1	0.06	55
4	Crack with moist argillaceous material	-45°	8×10^{-3}	22	1	0.065	30
5	Thick crack with moist argillaceous material	-45°	5×10^{-3}	35	1	0.05	80
6	Dielectric bedrock	Variable (undulating)	5×10^{-4}	2.75	1	0.18	60
7	Elongate cavity with moist ferrous material (laterite)	Quasi-horizontal with local irregularities	7×10^{-3}	22	1.4	0.055	22

The velocity structure is exactly known, therefore, it is possible to determine the amount of structural information recoverable from the radargram. Depth migration is a suitable approach and the matGPR implementation of the Split-Step method of Stoffa et al., (1990) was used, which is augmented to account for the frequency dependence of the phase velocity based on the formulation of Bano (1996). The depth-migrated radargram is shown at the bottom of Fig. 5a. The location and geometry of all objects is generally recovered with success but not with equal accuracy. Specifically, the lower tips of cracks 2, 3, 4 and 5 are disturbed by artefacts which are, at least partially produced by the interaction of their respective wavefronts with the wavefront of object 6 (bedrock) and with the clutter. Moreover, the interference of the clutter can clearly be seen in the images of cracks 3, 4 and 5, although it does not obscure them. Analogous but smaller scale effects can be observed around the central mound of the bedrock and at the lower surface of object 7 (cavity). The sub-vertical left and right interfaces of object 7 do not reflect signals from overhead sources and cannot be imaged.

The synthetic radargram was decomposed into six scales with the second coarser ($j=2$) scale comprising 32 angular wedges; the relatively fine angular decomposition is necessary because the scatterers (objects) are tightly spaced. In order to visualize how curvelet coefficients are stored and how they encode information, Fig. 5b shows the complete set for scales $j=1$ to $j=4$. It is apparent that significant information about the amplitude and geometry of the wavefronts comprising the radargram exists only in the north and south quadrants of scales $j=3$ and $j=4$. It must be emphasized that analogous significant information exists in the same quadrants of the finer fifth scale ($j=5$) which is not shown because the figure would then be rather difficult to read. The coefficients of the sixth scale are not curvelets; they are wavelets (isotropic) and at the relevant frequencies/ wavenumbers represent low-amplitude random noise. The cosine coefficients are arranged in the north and east quadrants (upper-right diagonal) and for $j=3$ and $j=4$ are indexed from $l=1$ to $l=32$ with respect to the NW corner. The sine coefficients are arranged in the south and west quadrants and for $j=3$ and $j=4$ are indexed from $l=33$ to $l=64$. The coefficients representing the structural elements 1 – 7 of the radargram at these scales are also indicated in the figure.




Figure 5b. The curvelet coefficients generated by the DCT of the synthetic radargram of Fig. 5a; only the scales $j=1$ to $j=4$ are shown for the sake of readability. The cosine coefficients are arranged in the north and east quadrants (upper-right diagonal). The sine coefficients are arranged in the south and west quadrants (lower-left diagonal). The coefficients encoding the structural elements 1 – 7 of the synthetic radargram for scales $j=3$ and $j=4$ are also indicated.

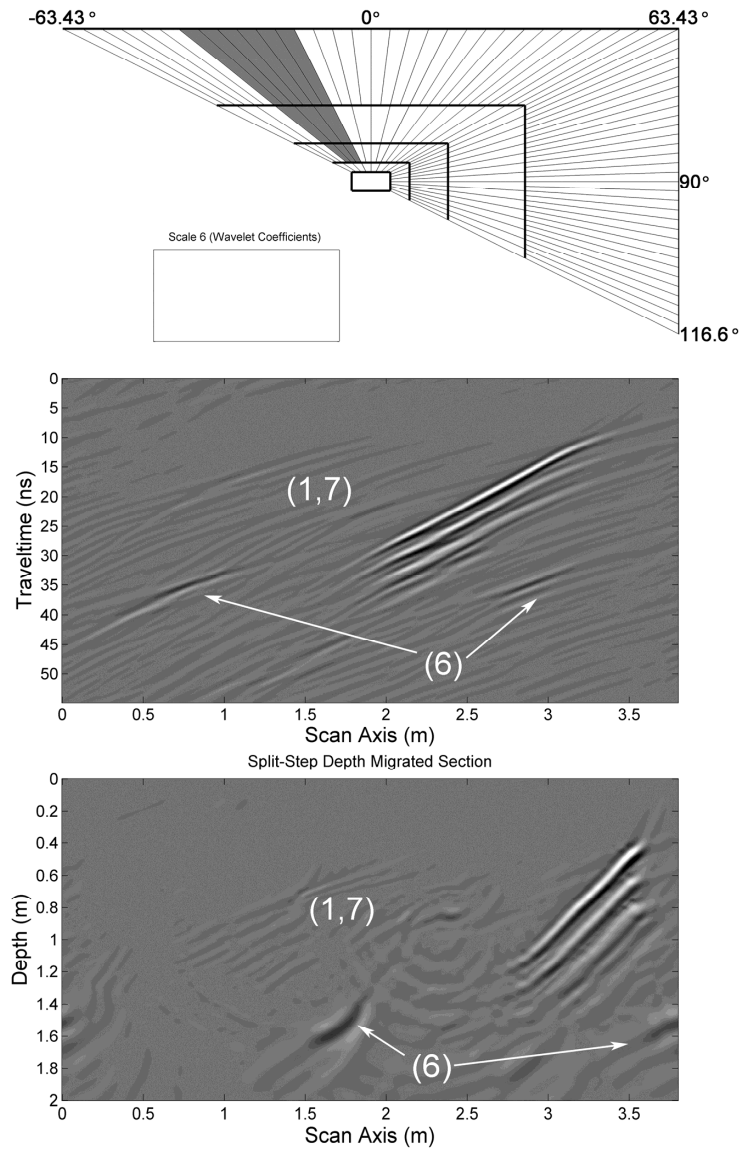


Figure 5c. Top: The selection of coefficients used for the extraction of cracks 3, 4 and 5 from the synthetic radargram of Fig. 5a; they are represented by the subset of trapezoidal wedges $\{j = 3, l \in [4, 6]\} \cup \{j = 4, l \in [4, 6]\} \cup \{j = 5, l \in [7, 12]\}$. **Middle:** Partial reconstruction of the synthetic radargram based on the selection of curvelet coefficients shown at the top; the reconstruction includes complete information about cracks 3, 4 and 5 only. **Bottom:** The partially reconstructed radargram after 2-D split-step depth migration.

Cracks 3, 4 and 5 have distinct orientational characteristics and they can be effectively isolated from the other structural elements of the radargram, as demonstrated in Fig. 5b and 5c. The top panel of Fig. 5c shows the selection of coefficients used for the partial reconstruction displayed in the middle: this is represented by the subset of trapezoidal wedges $\{j = 3, l \in [4, 6]\} \cup \{j = 4, l \in [4, 6]\} \cup \{j = 5, l \in [7, 12]\}$, with all the rest having been shrunk to zero. Although for the sake of simplicity the wedges are indexed only in the north and east quadrants, from $l=1$ to $l=L/2$, it is always important to bear in mind that each wedge represents the pair $(l, L/2 + l)$ of cosine and sine coefficients that is actually used in the reconstruction. It is immediately apparent that the main part of the wavefronts generated by cracks 3, 4 and 5 are very efficiently reproduced and that the optimal reconstruction property of the CT has effectively eliminated interference by wavefronts from the other structural elements of the model. This is also clearly observed in the migrated radargram shown at the bottom, in which the images of the cracks are sharper and better focused and the artefacts associated with their lower tips are significantly reduced.

However, they are not eliminated because the curvelets used in the analysis did not (and could not) eliminate the matching orientational attributes of the interference. Diffractions from the upper tips of the cracks have not been included in the reconstruction; evidently, this does not (and could not) affect the quality of the result. In addition, there are very few matching orientational characteristics between cracks 3, 4 and 5 and the other structural elements, which could be encoded in the coefficients and appear in the reconstructed image: these are clearly marked and include short segments of object 6 and ghosts of objects 1 and 7.

In the next example, the target of the analysis is crack 2, which shares some orientational attributes with the bedrock (object 6). The middle panel of Fig. 5d shows a partial reconstruction based on the coefficients represented by the wedges shown at the top; these are $\{j = 3, l \in [10, 12]\} \cup \{j = 4, l \in [10, 12]\} \cup \{j = 5, l \in [19, 24]\}$ with all the rest having been shrunk to zero. The wavefront generated by crack 2 is efficiently reproduced and is free of any interference caused by the other structural elements and the clutter; the artefact associated with its lower tip is also significantly suppressed (bottom). The matching down-dipping orientational attributes of the bedrock are also encoded in the coefficients and included in the reconstruction, as indicated in the middle and bottom panels respectively. For the same reason, one can also observe a ghost of object 7 and, at the top right of the reconstructed and depth migrated radargrams, ghosts of the upper tips of cracks 3 and 4; the latter are due to the down-dipping parts of the diffraction fronts associated with the tips of these cracks and are clearly visible at the extreme right of the reconstructed radargram (middle).

It is significant to point out that the forward and inverse CT is very stable and does not generate artefacts. This is easy to verify by experiment and can also be seen in Fig. 3a. The right panel of that figure illustrates four curvelets generated with the wrapping algorithm from the corresponding frequency domain constructs illustrated at the left panel: if there was any instability, e.g. Gibbs effects, it would show up in the x -domain representation of the right panel! Moreover, the optimal reconstruction property of the CT also guarantees that the reconstruction error is very small. With double precision arithmetic in a 64-bit machine, the peak signal-to-noise ratio (PSNR) associated with whole image reconstruction of synthetic and real GPR data, (i.e. forward and inverse decomposition without any shrinkage), is typically higher than 700dB. In the example of Fig. 5, after negating the 1st, 2nd and 6th scales which contain lower power contributions and small scale transients, the PSNR is still as high as 154dB. Such performance is typical!

Although the CT is stable, there is a certain risk of spurious low-amplitude oscillations when the data contain transients that are *incompletely* reconstructed. Fig. 5e presents an example. The bottom panel shows a partial reconstruction of the synthetic radargram after negating the subset of coefficients $\{j = 1\} \cup \{j = 2, l \in [1, 16]\} \cup \{j = 5, l \in [7, 12]\}$ shown at the top. This is equivalent to high passing the f - k spectrum above the ceiling of the 2nd scale (90 MHz and 2.9 m⁻¹ respectively), and narrow-band stopping the f - k spectrum between 780 – 1545 MHz and 22 – 45 m⁻¹ (boundaries of the 5th scale) as well as between the slopes -55° to -29° (dips of Cracks 3, 4 and 5). Because Crack 3 (and in some measure crack 4) is very thin, it does not produce significant attenuation and is associated with a sharp wavefront (transient) and a broadband f - k spectrum. The narrow-band stopping operation at the fifth scale (top) generates the oscillatory pattern indicated by the black arrow in the partially reconstructed radargram (bottom). Conversely, such oscillations are not present in the partial reconstructions of Fig. 5c and 5d because they have been generated with coefficient subsets that are broad-band in frequency and wavenumber, albeit narrow-band in angular content. Although the spurious oscillation is seldom a real issue, it is always advisable to exercise caution when using the CT filter on broadband (e.g. transient) features.

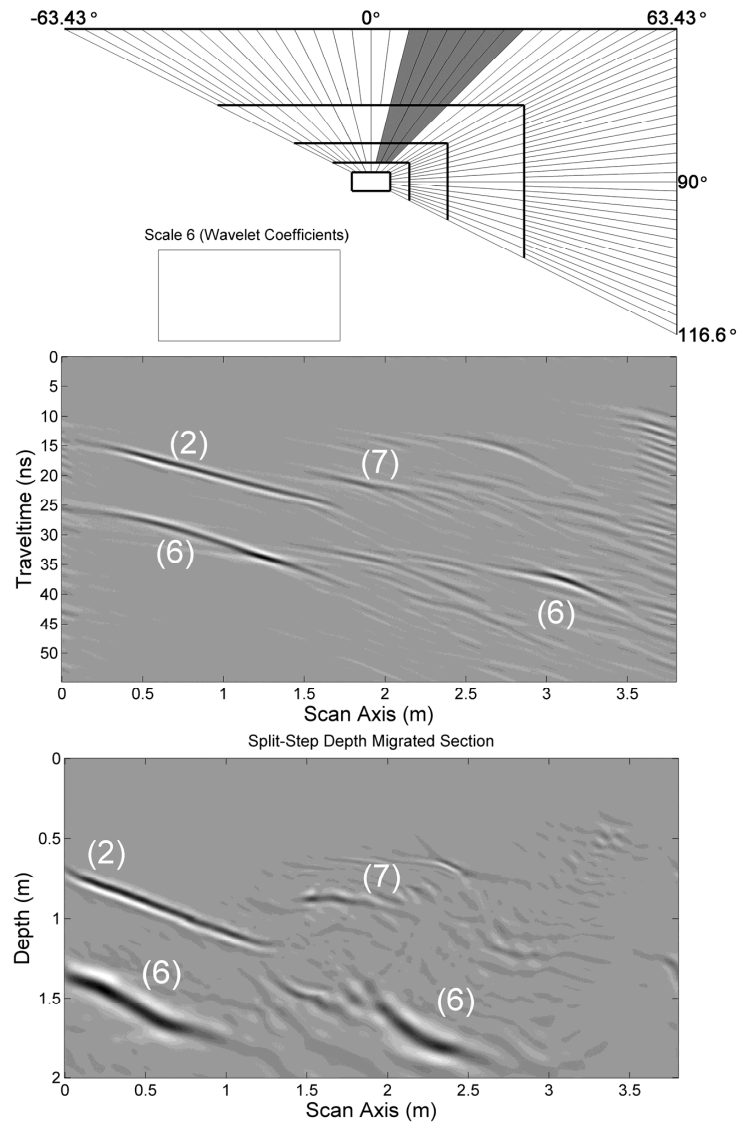


Figure 5d. Top: The coefficients used for isolating information about crack 2 from the synthetic radargram of Fig. 5a; they are represented by the subset of trapezoidal wedges $\{j = 3, l \in [10, 12]\} \cup \{j = 4, l \in [10, 12]\} \cup \{j = 5, l \in [19, 24]\}$. **Middle:** Partial reconstruction of the synthetic radargram based on the selection of curvelet coefficients shown at the top; the reconstruction includes complete information about crack 2, partial information about the bedrock (object 6) and a ghost of object 7 (cavity). **Bottom:** Split-step depth migration of the reconstructed synthetic radargram.

It is also important to point out the risk of misconstruing for artefacts (spurious oscillation) certain elements of partially reconstructed data. In the example of Fig. 5, the CT filter was used to perform a highly anisotropic analysis by targeting specific reflectors at specific orientations: curvelets and wavefronts were matched only along the orientations specified by the curvelets so that the partially reconstructed radargram contains almost uniformly dipping features, either localized or distributed at different locations. In general, if several small-scale such features appear distributed across a partially reconstructed radargram, they may conceivably create a false impression of artefacts due to the filtering process. In most cases, the false artefacts are directionally filtered wavefronts, as discussed above, or subtle data components not immediately evident to the naked eye. Note also that such features exist in every data set, inclusive of synthetics, and can be made to stand out in other ways, as can clearly be seen in the migrated synthetic radargram of Fig. 5a (bottom).

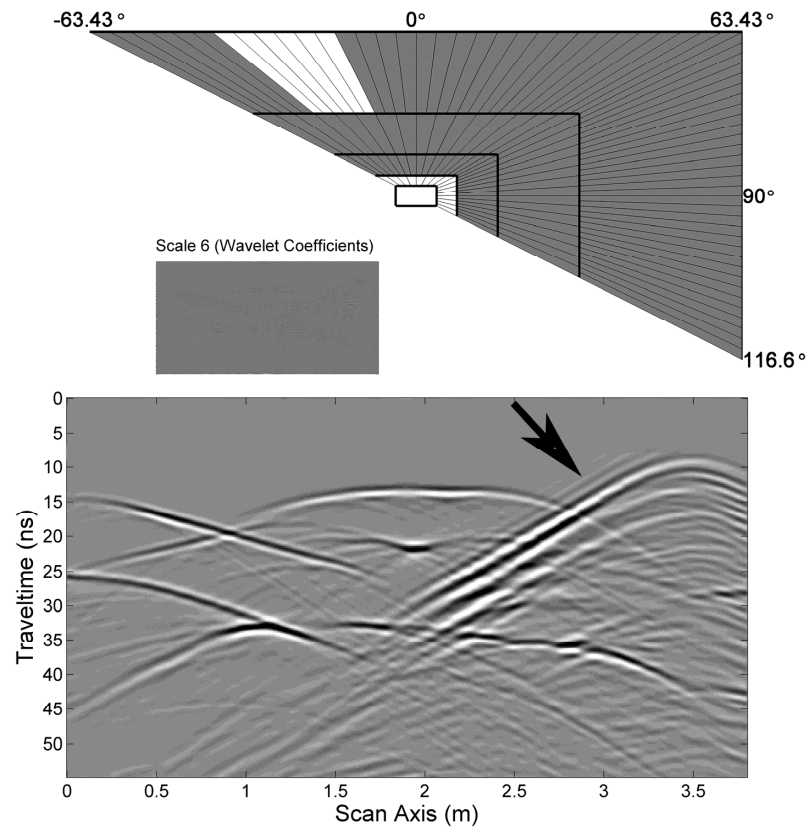


Figure 5e. Spurious oscillation due to narrow-band filtering of a transient wavefield. **Top:** The coefficients used in the partial reconstruction of the synthetic radargram of Fig. 5a. The subset $\{j = 1\} \cup \{j = 2, l \in [1, 16]\} \cup \{j = 5, l \in [7, 12]\}$ is *excluded* from the reconstruction. **Bottom:** Partial reconstruction of the synthetic radargram; the back arrow points to the low amplitude spurious oscillation generated by the narrow band-stopping operation.

3 Examples

The utility and versatility of the Curvelet Transform will now be demonstrated with five applications to data from archaeometric, geotechnical and hydrogeological surveys. The first three examples will demonstrate its usefulness in retrieving information from data contaminated by high levels of noise and featuring straight or curved reflections from internal interfaces. The fourth and fifth examples will demonstrate the effectiveness of the CT to process normal (average quality) data sets collected in complex propagation media. The first, second and fifth data sets have also been presented by Tzanis (2013), as part of the demonstration of directional and multidirectional wavelet filters. In this way, comparison between effective, single-scale directional filtering methods and the multi-scale Curvelet Transform will also be possible. For conciseness, only essential information about these example data sets and the context/environment in which they have been collected will be given herein; the interested reader may find additional details in the cited reference. Additional applications of the CT and comparisons with the multidirectional wavelet filters can also be found in Tzanis (2014).

The data of Fig. 6a was collected as part of an archaeometric survey with a GSSI SIR-2000 system and an antenna with a nominal central frequency of 400MHz. The raw radargram is shown as measured: it is quite noisy and comprises a 512 sample \times 1024 traces section with a sampling rate of 0.1957ns (total time window = 100ns) and trace spacing 0.01075m (section length = 11m). Two apparent and significant

features in this section are the signature of a buried wall¹ at distances 1.5 – 2.5 m and traveltimes 30 – 70 ns and an up-dipping reflector which is clearly seen between coordinates (60ns, 6m) and (49ns, 7.8m), while there's quite clear indication that it may extend bilaterally to later times/shorter distances (approx. 66ns/5m) and earlier times/longer distances (approx. 37ns/10m).

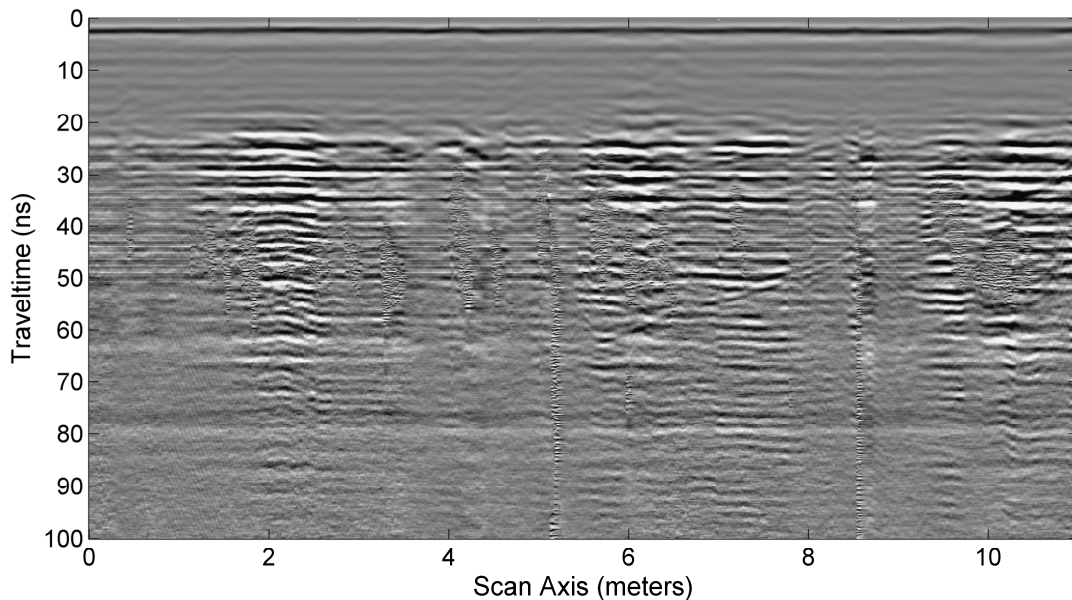


Figure 6a. B-scan radargram featuring the signature of a buried wall between the ordinates 1.5 – 2.5 m and traveltimes 30–70 ns, as well as a linear up-dipping reflector between the ordinates 49 – 60 ns and abscissae 6 – 7.8 m. The linear reflector is deeply buried in noise.

The radargram was decomposed into six scales with the second coarser ($j=2$) scale comprising 28 angular wedges. After some straightforward experimentation, it is easy to see that the main noise components comprise: a) Very low frequency isotropic interference ($j = 1, f < 110$ MHz, $k < 1.9$ m⁻¹); b) Low frequency, mainly horizontal ringing with frequencies $f \leq 210$ MHz. c) High frequency bursts, localized and mainly horizontal, such that $f \geq 848$ MHz; d) Broadband spatial variation of vertical to sub-vertical orientation. The noise can be almost precisely excised by negating all the curvelet coefficients save for those shown in Fig. 6b (left), i.e. by partially reconstructing the data based on the subset of coefficients $\{j = 3, l \in [1, 14] - [7, 8]\} \cup \{j = 4, l \in [1, 14]\}$. The image of the wall at around 2 m is now clear. Images of additional, possibly man-made structures are also apparent between the ordinates 6 m – 9 m and 25 ns – 50 ns, while the up-dipping reflector is also more clearly discerned. There is some residual ringing left, which cannot be removed without detrimental effects to the data – it is represented by the curvelets $\{j = 4, l \in [7, 8]\}$. The up-dipping reflector can be isolated by negating all the curvelet coefficients except for those in $\{j=4, l \in [4, 6]\}$, as shown in Fig. 6c. Notably, this subset comprises curvelets exactly parallel ($j = 5$) and sub-parallel to the dip of the reflector and frequencies 429 MHz – 848 MHz. The dipping reflector, not only stands out clearly and its lateral extent beyond the initially observable range is confirmed, but is also optimally recovered and appears to be almost continuous across the section owing to the special microlocal features which render curvelets ideal for reconstruction problems with missing or hidden data (what you see is what you *can* get).

¹ This is an interpretation. The same reflection pattern extends laterally to either side of this particular radargram. The survey area abounds with buried walls – it is adjacent to the Agora (Forum) of the Classical and Roman Argos (Greece). Such buried walls have been verified by GPR (same type of signature) and excavation, the nearest trench being approx. 25m away.

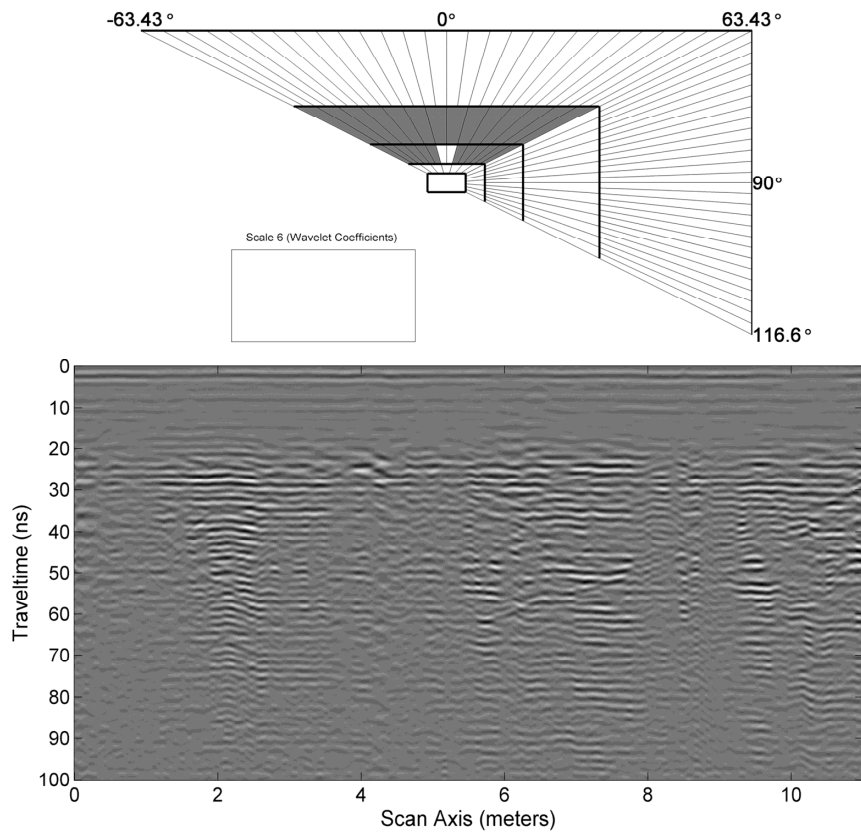


Figure 6b. The bottom panel illustrates the de-noised radargram of Fig. 5a, after negating “noise” coefficients as shown in the top panel.

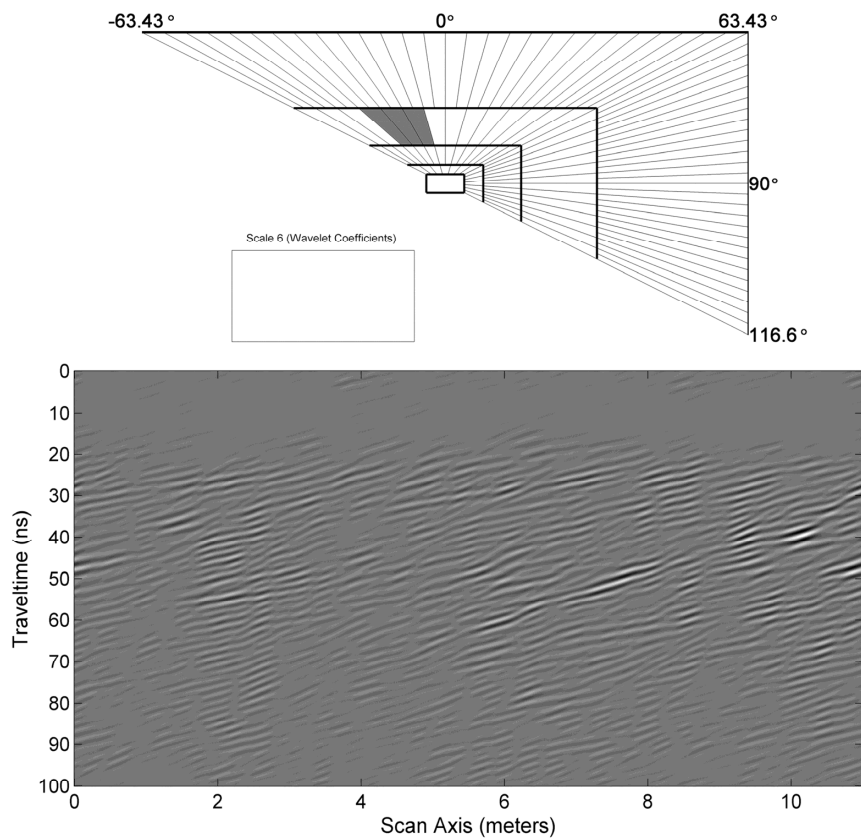


Figure 6c. The bottom panel illustrates the up-dipping reflector featured in Fig. 5a and Fig5b, reconstructed from the curvelet coefficients $\{j=4, l \in [4, 6]\}$ shown in the top panel.

The second example may be familiar to several GPR practitioners: it is the radargram distributed with the GPR analysis package of Lucius and Powers (2002). The original section was measured in B-scan equal time spacing mode at the Norman (Oklahoma) Landfill with a GSSI SIR-2000 system and a 500MHz low power antenna; it was subsequently pre-processed (transformed to equal trace spacing and resampled to 512 samples \times 512 traces, so that it now has a sampling rate of 0.1988 ns (time window equals to 101.7ns) and a trace spacing of 0.0387 m (section length is 19.8m). The section is shown in Fig. 7a and can be seen to suffer from crossing clutter, characteristic of multiple small targets or rough reflective surfaces. The noise can be locally strong, but it does not completely overshadow the data which is still interpretable. Accordingly, in this example the performance of the proposed analysis scheme can be precisely evaluated because the observer can see exactly what lies behind the noise.




Figure 7a. The radargram distributed with the GPR analysis package of Lucius and Powers (2002), transformed to equal trace spacing and resampled to a 512 \times 512 matrix. The data suffer from crossing clutter, characteristic of multiple small targets or rough reflective surfaces.

In this example, the radargram was decomposed into six scales with the second coarser ($j=2$) scale comprising 24 angular wedges (the number of wedges doubles in every second scale). The clutter waveforms dip at high angles and comprise spatial rather than temporal features. They comprise two groups, one with shorter spatial widths and higher intensity, as for instance between the ordinates 2 – 4 m and traveltimes 20 – 60 ns, and one with longer spatial widths and lower intensity, as for instance between the ordinates 9 – 12 m and traveltimes 60 – 80 ns. The spatial widths of both groups can be roughly measured at several locations in the radargram; they average to approx. 0.25m, which would imply expected wavenumber(s) of the order of 4m^{-1} and would place them near the boundary of the fourth ($2.18\text{ m}^{-1} < k < 4.29\text{ m}^{-1}$) and the fifth ($4.29\text{ m}^{-1} < k < 8.63\text{ m}^{-1}$) scales. Fig. 7b illustrates a model of the clutter (right panel) reconstructed from the curvelet subset $\{j = 4, l \in [13, 24]\} \cup \{j = 5, l \in [25, 48]\}$, as shown in the left panel.

A simple inspection of the data will show that the main reflections from subsurface are interfaces exhibit apparent dip shallower than 45° and are associated with spatial widths (scales) of metric order. Likewise, an inspection of individual trace spectra, as well as of the f-k spectrum will show that the data is disproportionately rich in low frequencies, with the peak located in the neighbourhood of 300MHz (on average). This would place the main components of the structural information in the third and fourth

scales where $207 \text{ MHz} < f < 835 \text{ MHz}$ and $1 \text{ m}^{-1} < k < 4.29 \text{ m}^{-1}$. Thus, it is possible to obtain a representation of the data without clutter and low frequency/short wavenumber interference and without significant loss of structural information. This can be achieved via a partial reconstruction based on the curvelet subset $\{j = 3\} \cup \{j = 4, l \in [1, 12]\} \cup \{j = 5, l \in [1, 24]\}$; the result is shown in Fig. 7c.

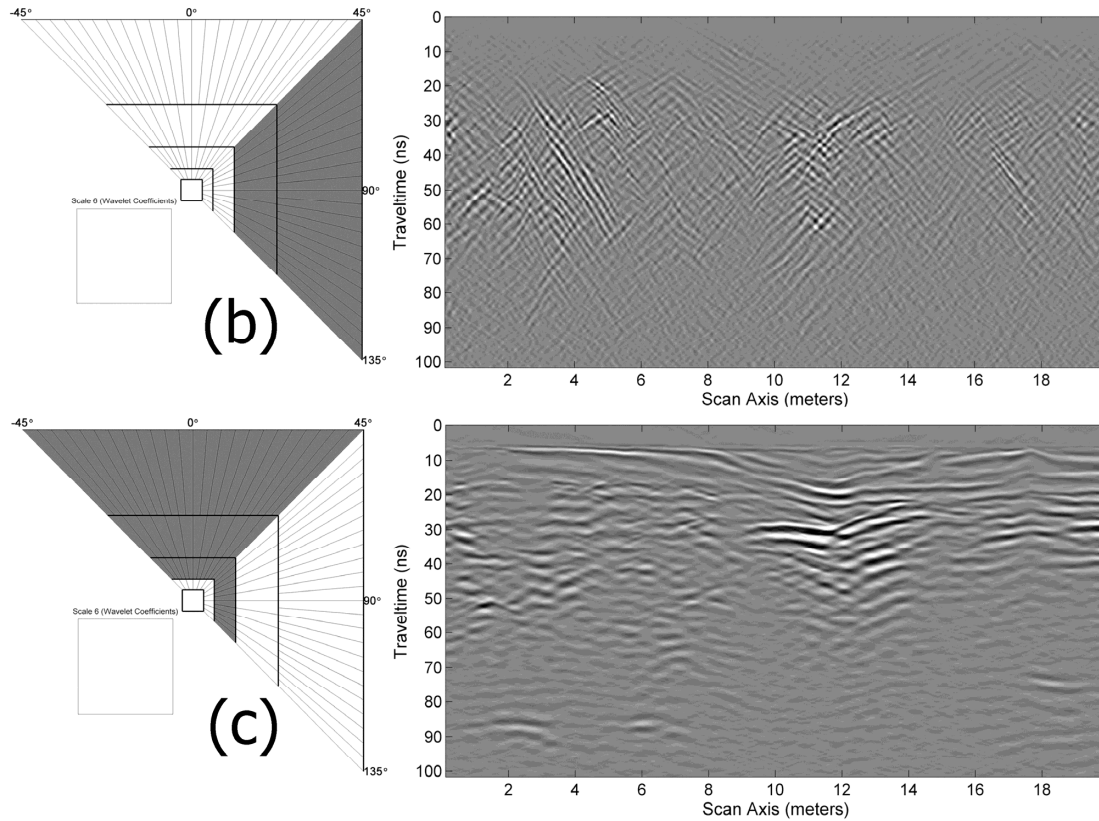


Figure 7. (b) The top-right panel illustrates a partial reconstruction of the data of Fig. 7a based on curvelets in the subset $\{j = 4, l \in [13, 24]\} \cup \{j = 5, l \in [25, 48]\}$ shown in the top-left panel. This effectively isolates the crossing noise process. (c) The bottom-right panel illustrates a clutter-free partial reconstruction of the data of Fig. 7a based on curvelets in the subset $\{j = 3\} \cup \{j = 4, l \in [1, 12]\} \cup \{j = 5, l \in [1, 24]\}$, as shown in the bottom-left panel.

The third example will demonstrate the recovery of structural information from very noisy data measured in an adverse geological setting. The data was collected at the Kato Souli plain, near the east margin of the Marathon basin (NE Attica, Greece) and in fallow field located at the north border of the Schinias wetland (natural reserve). The section begins at coordinates (24.034825°E, 38.169714°N) and has a length of 141.8m and an azimuth of N162° – it can be very easily located with free on-line mapping software, for instance GoogleEarth™. The ground consists of thick, generally moist Holocene alluvial sediments with a significant argillaceous component and contains a shallow (< 2m) unconfined aquifer with brackish water, as it appears to be recharged from deeper water beds that are expressly saline due to extensive sea water intrusion. The data shown herein was collected with a Måla GPR system and 250MHz antenna, as part of an experiment to evaluate the feasibility of monitoring seasonal water table variations in a high attenuation environment. Fig. 8a shows the example section after time-zero adjustment, global background removal, time gain with the “inverse amplitude decay” method and resampling to a 256 sample \times 512 trace matrix, so that the final sampling rate is 0.4848 ns and the trace spacing 0.2756 m. Twenty metres to the NNW from the beginning of the section and almost in-line with

the scan axis, there is a water well in which the level of the water table could be measured at approx. 1.6m below ground level and the salinity was found to be 50 mgr/lit in Na^+ . The end of the section is approx. 30 m away from the waterfront of the Schinias marsh. The difference in elevation between the well and the waterfront is approx. 1.5 m.




Figure 8. (a) Radargram obtained at the margin of the Schinias marsh, NE Attica, Greece. The data has been pre-processed with time-zero adjustment, global background removal and time-dependent amplification (time gain). **(b)** The data of Fig. 8a after high-pass Karhunen-Loeve filtering (eigenimages $p=3$ to $q=256$) and static correction with a velocity of 0.085 m/ns.

The data can be seen to suffer from intense, (sub)horizontal antenna self-clutter and severe, long wavenumber crossing clutter characteristic of multiple small targets and/or rough reflective surfaces. There's also intense (sub)vertical spatial variation due to corresponding discontinuities in the pattern of the horizontal self-clutter, as well as the signature of an (identified) shallow-buried small metallic object at approx. 115 m down the scan axis. The noise is generally overwhelming and the only part of the radargram with hints of earth structural information is a strip between 10 ns and 40 ns along the section; this is expected to contain the response of the water table.

The intensity of the self-clutter and redundant, small scale data components making up the high frequency/long wavenumber noise structure may be reduced, so as to alleviate the burden of complexity to be treated with the curvelet transform. Herein, this is done with the *Karhunen-Loeve Transform* (KLT -

e.g. Fukunaga, 1990), also known as the *Eigenimage* or *Eigenvector Transform*. The KLT performs the same function as the Fourier transform but with a different set of basis functions that can be derived from the (auto)covariance matrix of the process, depend on the particular matrix (image) and enable its decomposition in as an economical way as possible, thus yielding an expansion with coefficients that are truly uncorrelated. The KLT can be realized with the singular value decomposition, which for a rank R matrix \mathbf{X} with M rows and N columns is

$$\mathbf{X} = \mathbf{U} \cdot \mathbf{S} \cdot \mathbf{V}^T = \sum_{i=1}^R s_i \mathbf{u}_i \mathbf{v}_i^T .$$

\mathbf{S} is diagonal and contains the singular values of \mathbf{X} , namely the positive square roots of the eigenvalues of the covariance matrices $\mathbf{X}\mathbf{X}^T$ and $\mathbf{X}^T\mathbf{X}$, arranged in non-increasing order; s_i is the i -th singular value of \mathbf{X} . \mathbf{U} is M -by- M orthogonal with \mathbf{u}_i being the i -th eigenvector of $\mathbf{X}\mathbf{X}^T$; \mathbf{V} is N -by- N orthogonal and \mathbf{v}_i the i -th eigenvector of $\mathbf{X}^T\mathbf{X}$, so that and $\mathbf{u} \mathbf{v}^T$ is an M -by- N orthogonal matrix representing the i -th eigenimage of \mathbf{X} . In general, it is possible to reconstruct a partial representation of \mathbf{X} from only some eigenimages, i.e.

$$\tilde{\mathbf{X}} \approx \sum_{i=p}^q s_i \mathbf{u}_i \mathbf{v}_i^T, \quad 1 < p \leq q < R. \quad (11)$$

This is tantamount to band-pass filtering the matrix in the energy (size) scale $[p, q]$. Low-pass and high-pass filtering is also possible by performing the summation from $i=1$ to $i=p$ and from $i=q$ to $i=R$ respectively. The application of the KLT in the analysis of GPR data is not very common. It has been extensively discussed by Cagnoli and Ulrych (2001), while there are few examples in conference presentations (e.g. Zhao et al. 2005; Rudzki, 2008; Xie et al. 2013; a very few others).

As is apparent in Fig. 8b, the self-clutter is a large scale process – it is straightforward to verify that the eigenimages associated with the first two singular values of the data matrix are exclusively associated with it. In consequence, it is possible to derive a partial (high-pass) reconstruction of the original data set, by setting $p = 3$ and $q = 256$ in eq. 11, so that the self-clutter is partially suppressed *without* loss of structural information. Given also that precise levelling measurements were conducted along the section, it was possible to apply a static correction, using the elevation of the water well as reference and a velocity of 0.085 m/ns, obtained as the average of a set of direct measurements in the field. The final (partially de-noised and statically compensated) section is shown in Fig. 8b; it is apparent that the high intensity/ low frequency/ short wavenumber noise has been considerably reduced.

The Curvelet Decomposition applied to the data of Fig. 8b comprised of 6 scales and 28 angles at the second coarser scale; the short size of the vertical axis (256 samples) and relatively long sampling interval do not allow higher order decomposition. Quite apparently, the self-clutter is rather broadband and is represented by (sub)vertically oriented curvelet coefficients. The remaining part of the noise is principally of high apparent dip – broadband and expressly manifested in scale 2 (low frequency/ short wavenumber, $0.078\text{m}^{-1} < k < 0.149\text{m}^{-1}$ and $40\text{MHz} < f < 89\text{MHz}$), in scale 4 ($0.305\text{m}^{-1} < k < 0.602\text{m}^{-1}$ and $169\text{MHz} < f < 346\text{MHz}$, around the nominal central frequency of the antenna), as well as in scales 5 and 6 (high frequency/ long wavenumber). After some experimentation, the reconstruction was based on the subset $\{j = 2, l \in [1, 3] \cup [5, 9]\} \cup \{j = 3, l \in [1, 6] \cup [10, 18], \cup [25, 28]\} \cup \{j = 4, l \in [1, 6] \cup [10, 18] \cup [25, 28]\} \cup \{j = 5, l \in [1, 12] \cup [17, 28]\}$, as shown in Fig. 8c-top. The reconstructed data is shown in Fig. 8c-bottom. One may note that the reconstruction is not exactly noise-free: residual crossing clutter is apparent but is severely attenuated so as not to interfere with the structural data. The strip between 10 ns and 40 ns now exhibits definite evidence of a laterally extended, quasi continuous structure, which contains the response of the water table. The response is rather irregular, possibly due to a correspondingly irregular soil – water interface and also due to the residual noise.

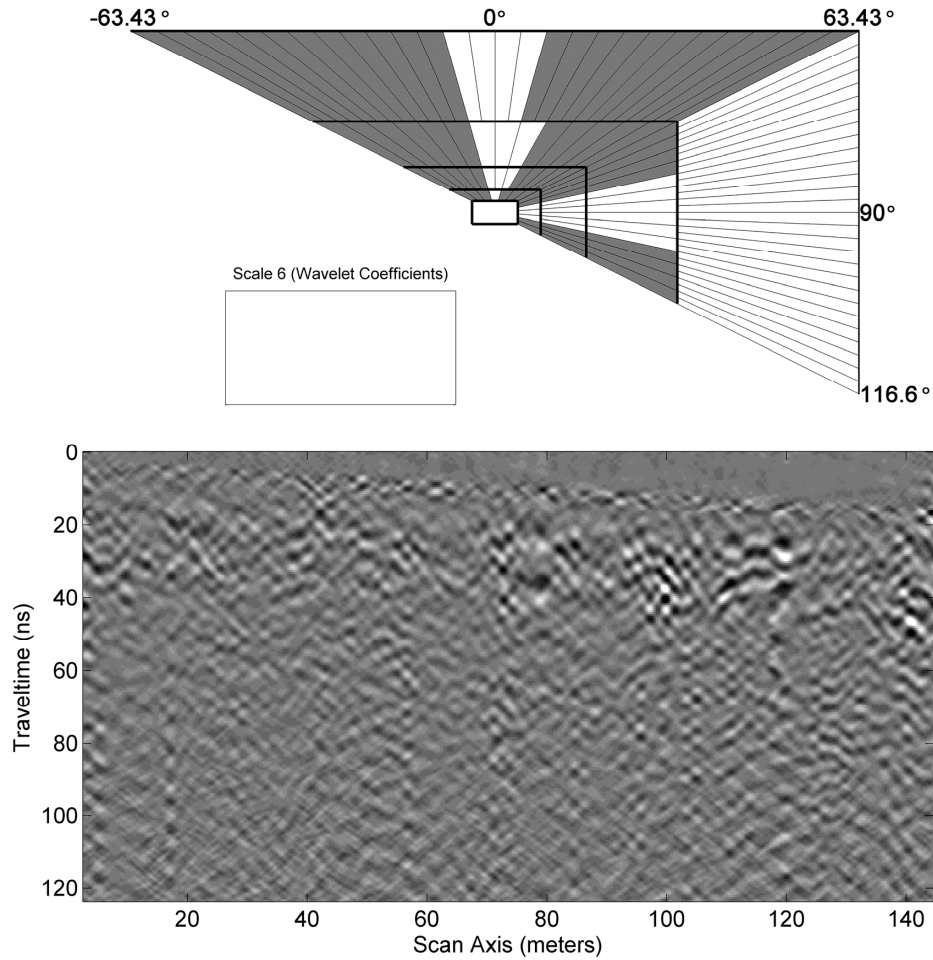


Figure 8c. The bottom panel illustrates a partial reconstruction of the section shown of Fig. 8b, based on the subset of curvelet coefficients $\{j = 2, l \in [1, 3] \cup [5, 9]\} \cup \{j = 3, l \in [1, 6] \cup [10, 18], \cup [25, 28]\} \cup \{j = 4, l \in [1, 6] \cup [10, 18] \cup [25, 28]\} \cup \{j = 5, l \in [1, 12] \cup [17, 28]\}$ shown in the top.

In order to demonstrate that this is indeed the response of the water table, the depth-migrated traces of the median instantaneous amplitude and mean centroid frequency are compared in Fig. 8d. The median instantaneous amplitude is the median of the instantaneous amplitudes computed from the analytic signals of all traces of the reconstructed radargram. The centroid frequency is the frequency of the “centre of mass” of the signal spectrum (spectral centroid) and provides a synopsis of the spectral content, hence a measure of changes in propagation conditions. It is calculated as a weighted mean of the frequencies present in the signal:

$$f_c = \frac{\int_0^{\infty} fS(f)df}{\int_0^{\infty} S(f)df},$$

where $S(f)$ is the amplitude spectrum. When the amplitude of high frequency spectral content is preferentially shrinking due to the signal propagating in high attenuation domains, the centroid frequency is expected to shift to lower frequencies. Conversely, it is expected to shift to higher frequencies when the signal enters low attenuation domains, or when it is sufficiently weak for high frequency spectral components and random noise to prevail. It is also expected to exhibit persistent gradual downshift in cases of dispersive propagation (e.g. Irving and Knight, 2003). Herein, for each trace of the reconstructed data, the corresponding trace of the centroid frequency was computed on the basis of an ultra-high

resolution time-frequency representation enabled by the S-Transform (Stockwell et al., 1996). Respectively, the mean centroid frequency trace is the average of all centroid frequency traces computed from all traces of the reconstructed data.

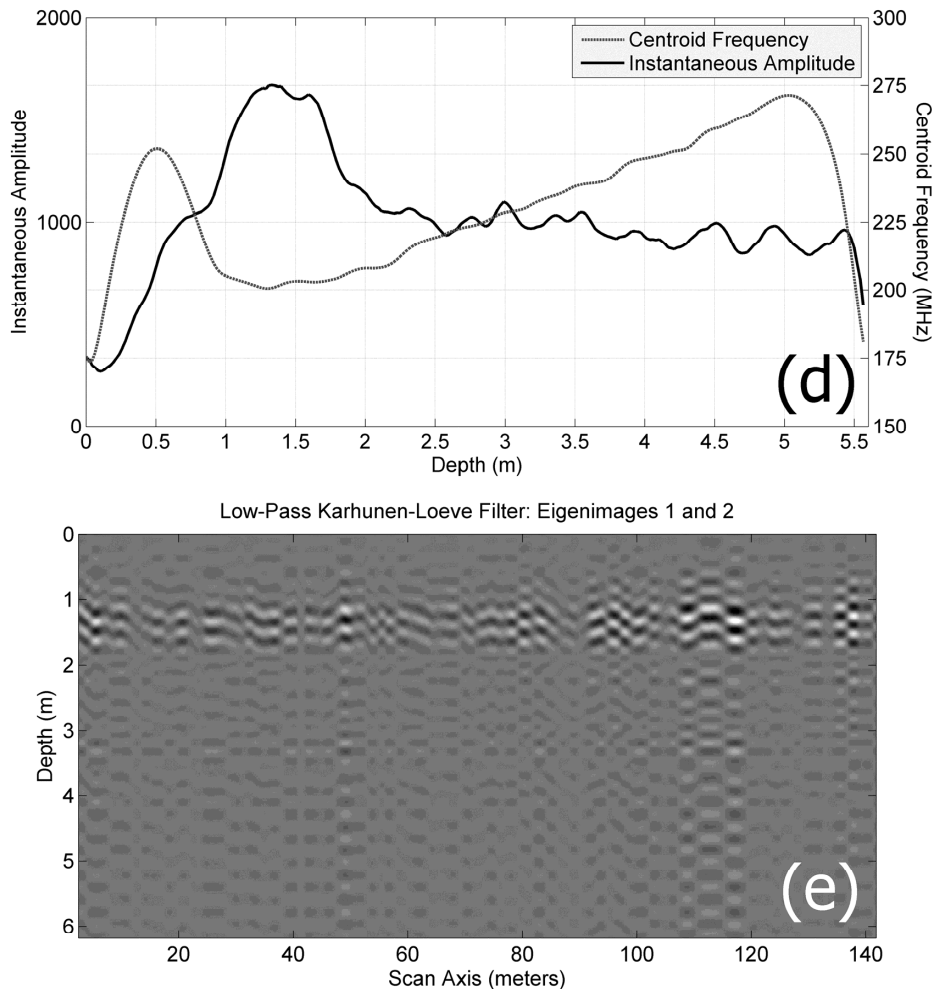


Figure 8. (d) Depth-migrated traces of the median instantaneous amplitude and mean centroid frequency computed from the filtered data shown in Fig. 8c. (e) The filtered data of Fig. 8c after low-pass Karhunen-Loeve filtering (eigenimages 1 and 2).

On studying Fig. 8d, it is immediately apparent that in the depth range 0.8 – 2.5 m, the median instantaneous amplitude has a pronounced broad peak that maximizes in the interval 1.3 – 1.7 m. The mean centroid frequency has a valley in the same depth range, exhibiting rapid decrease in the interval 0.5 – 1 m and dropping to approx. 200MHz in the interval 1.3 – 1.7 m before gradually increasing to approx. 270MHz at the depth of 5m. The opposite variation of the two attributes indicates the dominant presence of reflected low-frequency spectral components and the conspicuous absence of high frequency spectral components in the interval 1.3 – 1.7 m. Consistently with the theoretical and experimental work of Bano (2006, 2007) this is characteristic of a high attenuation domain, which in our case is the salinized unconfined water table and its capillary fringe. Also consistently with Bano (2006, 2007), the interval 0.5 – 1m, where precipitous attenuation of high frequencies is observed, is possibly related to the transition between the unsaturated and saturated zones. Although the soil – water interface is irregular by nature (formation) and complexity of the capillary fringe, the water table response is the largest scale structure in the filtered data set of Fig. 8c. Accordingly, it can be isolated with the aid of the KLT. The top panel of

Fig. 8d shows a depth-migrated reconstruction of the curvelet filtered radargram, based on the first two eigenimages only! This is an almost optimally recovered response of the water table, which is observed at a constant level of approx. 1.3 – 1.6m below the reference elevation (water well) and approaches the surface towards the end of the section (towards the waterfront); the range of depths at which the water table is detected, is consistent with the control points at the beginning (water well) and end of the section (waterfront).

The fourth example uses a “normal” data set with relatively high S/N ratio in order to demonstrate the capability of the curvelet transform to extract geometrical information at any scale or combination of scales. The survey was conducted on a levelled surface above massive limestone fragmented by (now inactive) conjugate normal faulting and jointing; karstification has nucleated along the fault walls and many of the resulting voids were subsequently filled with ferriferous argillaceous material (Fig 9a). The top panel of Fig. 9b illustrates a B-scan radargram collected with a Måla GPR system and 250MHz antenna. The data is shown after pre-processing with time-zero adjustment, global background removal, amplification with the “inverse amplitude decay” technique (e.g. Tzanis, 2010, p25) and resampling to a 1024 sample \times 1024 trace matrix (the sampling rate is 0.1909 ns and trace spacing 0.0313 m).



Figure 9a. Close-up snapshot of the fragmented limestone setting in which the data of Example 4 has been recorded. Void or laterite-filled faults and joints can clearly be observed. Photographs are courtesy of Mr P. Sotiropoulos, Terra-Marine Ltd., Greece (<http://terra-marine.gr>).

Noise is not a significant problem for this data, in which the main structural features are:

- a) Down-dipping reflections attributed to faulting. The most prominent such feature can be observed between 20 to 180ns and 2 to 20m. Along the trace of this structure one may observe clusters of strong reflections corresponding to areas of intense fragmentation and karstification. This process has led to the formation of larger scale structures, such as is the lens-shaped one between the ordinates 12 – 20 m and traveltimes 100 – 150 ns.
- b) Very faint steeply up-dipping linear reflections due to short, small-aperture antithetic fractures. There appears to be an association between the clustering of these fractures and the clustering of reflections along the fault walls.

The linear up-dipping reflections can be clearly observed in the *most positive curvature* attribute of the radargram; the most positive curvature (e.g. Roberts, 2001) is a very effective edge locator and is illustrated at the bottom panel of Fig. 9b. The most significant up-dipping reflections are indicated with

black arrows. The down-dipping reflections can also be clearly observed in the maximum curvature image, but due to their origin in the anomalous fault zone, they are evidently more complex and not as highly localized as the up-dipping features. Since the detection of fractures has always been a major application of GPR, this data is suitable for demonstrating the utility of the curvelet transform.

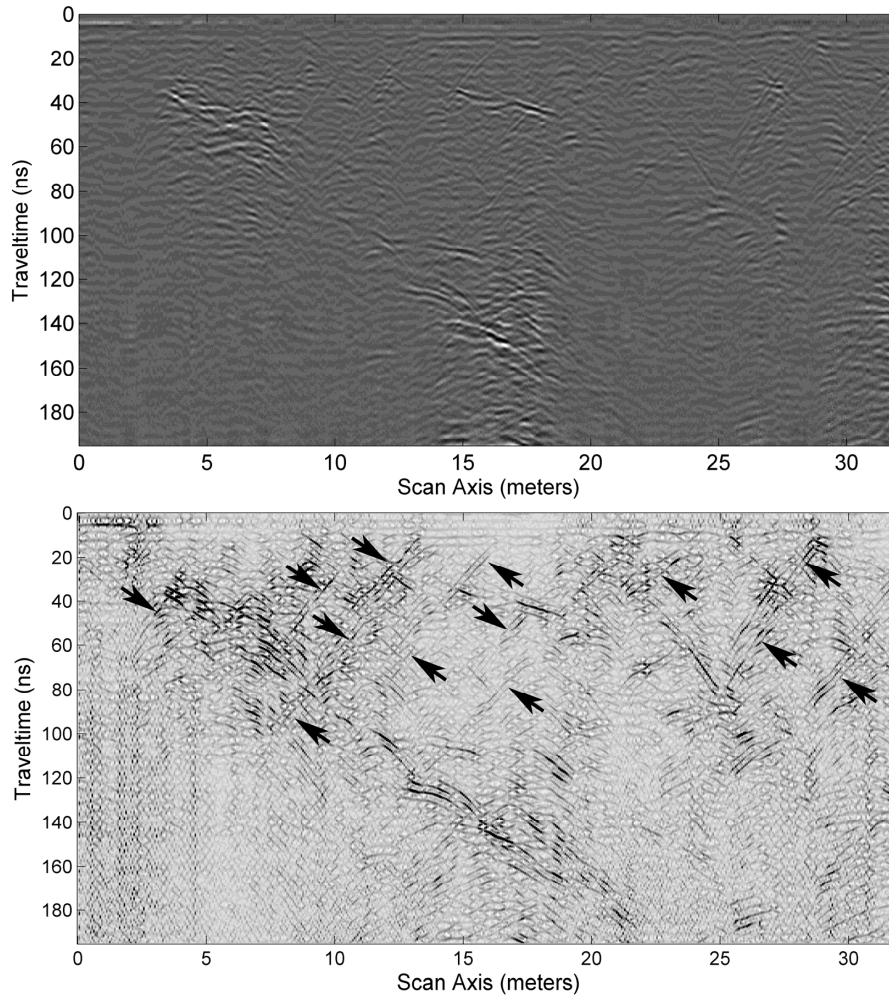


Figure 9b. Top: A B-scan radargram obtained above massive fragmented limestone. The data has been pre-processed as detailed in the text and is used courtesy of Mr P. Sotiropoulos, Terra-Marine Ltd (<http://terra-marine.gr>). **Bottom:** The most positive curvature attribute computed from the radargram above. The black arrows point to weak linear up-dipping reflections which are discernible but very faint in the radargram.

Because the steeply up-dipping linear reflections exhibit rather fine spatial and temporal scales and a limited range of dip angles, the Curvelet Decomposition comprised 7 scales, 32 angles at the second coarser ($j = 2$) scale. The up-dipping reflectors have spatial widths of the order of $0.2 - 0.3$ m, therefore expected wavenumbers of the order of $3.3 - 5$ m^{-1} . This means that for the matrix size and sampling rates of this data, they belong mainly to the 5th scale (2.65 $\text{m}^{-1} < k < 5.33$ m^{-1}). Their apparent dip (in data matrix coordinates) can also be measured and is found to range between -70° and -60° . Accordingly, the reconstruction considers only curvelets supported on wedges sub-normal to the apparent dip of the up-dipping reflections. Furthermore, it is based mainly on scale 5, but because several of these events are transient, it also includes parts of scales 4 and 6 so as to avoid spurious oscillation. The results are shown in Fig. 9c, which illustrates a partial reconstruction based on the coefficients $\{j = 4, l \in [27, 28]\} \cup \{j = 5, l \in [52, 57]\} \cup \{j = 6, l \in [53, 56]\}$. The black arrows pointing to the reconstructed up-dipping

reflections are placed exactly as per Fig. 9b; as scrutiny may show the reconstruction accurately represents one for one the images of the small-aperture fractures which can now be identified, mapped and analysed much easier (compare with analogous result of Fig. 5c). Notably, the low frequency/ short wavenumber noise, as well as all every geometrical feature outside the reconstruction range has disappeared from the processed radargram.

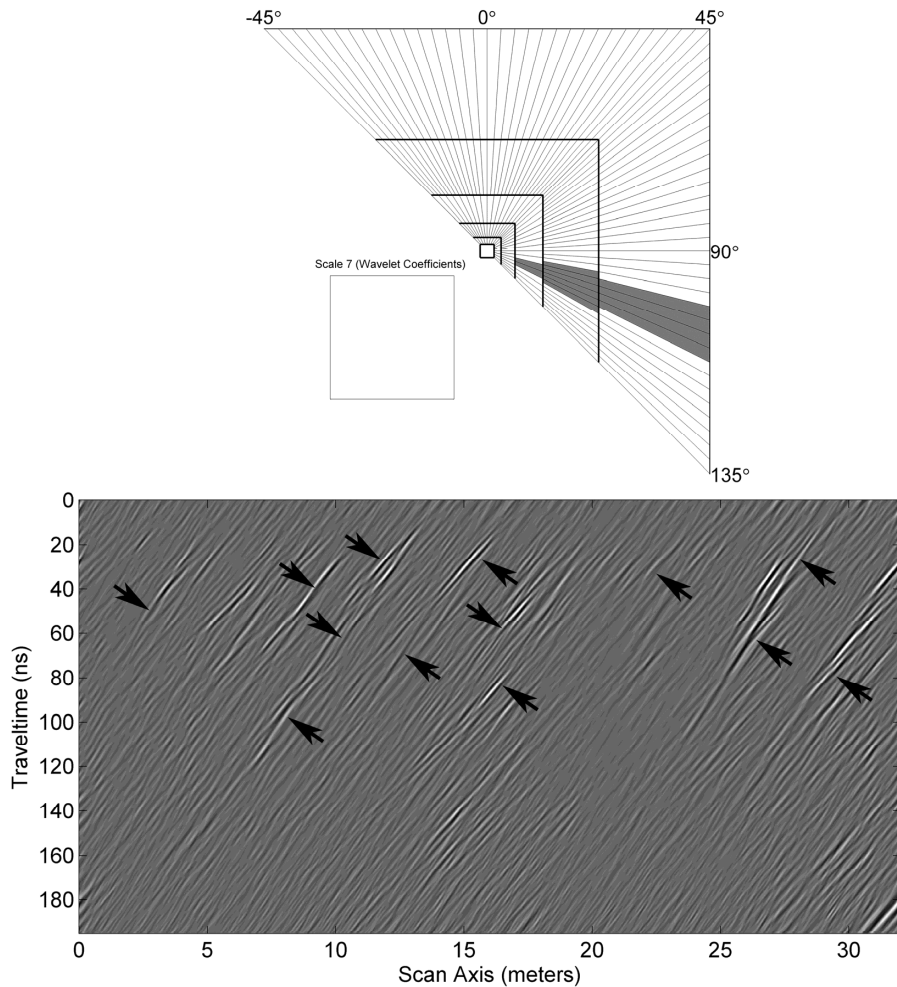


Figure 9c. The bottom panel illustrates a partial reconstruction of the radargram shown in Fig. 9b, based on the subset of coefficients $\{j = 4, l \in [27, 28]\} \cup \{j = 5, l \in [52, 57]\} \cup \{j = 6, l \in [53, 56]\}$ shown at the top. The reconstruction effectively isolates steeply up-dipping linear reflections from small aperture antithetic fractures (compare with analogous result of Fig. 5c). The black arrows pointing to the reflections are placed identically as in Fig. 9b (bottom).

The middle panel of Fig. 9d illustrates the up-dipping reflectors together with a reconstruction of the down-dipping reflections. The latter is based only on scales 4 and 5 because the down-dipping reflections are coarser and their high frequency/ long wavenumber spectral content is negligible. It is also based on wedges with slopes $50^\circ\text{--}74^\circ$, which are sub-normal to their apparent dips (in matrix coordinates). Accordingly, as shown in the top panel of Fig. 9d, the reconstruction of the down-dipping reflections included the coefficients $\{j = 4, l \in [18, 22]\} \cup \{j = 5, l \in [35, 44]\}$. With the interference from wavefronts with different geometries eliminated, the configuration of the synthetic (main) fractures can be seen to stand out rather clearly. Moreover, the relative significance of the synthetic and antithetic fractures, as indicated by their different reflectivities, is easier to evaluate. The bottom panel of Fig. 9d

illustrates the most positive curvature attribute of a time and depth migrated version of the section shown in the middle, plotted so that the data units are the same in both directions (1:1 aspect ratio). The migration is only approximate and based on a uniform velocity of 0.085 m/ns, estimated by fitting diffraction front hyperbolae observed elsewhere in the study area. Nevertheless, the up- and down-dipping reflections form a pattern such that the acute angle between them is approx. 60° . Given the geological setting and history of the area, this is additional evidence of conjugate synthetic and antithetic faulting and jointing. Overall, Fig. 9d provides a much clearer picture of the local fragmentation pattern, than what was afforded by the original low-noise radargram.

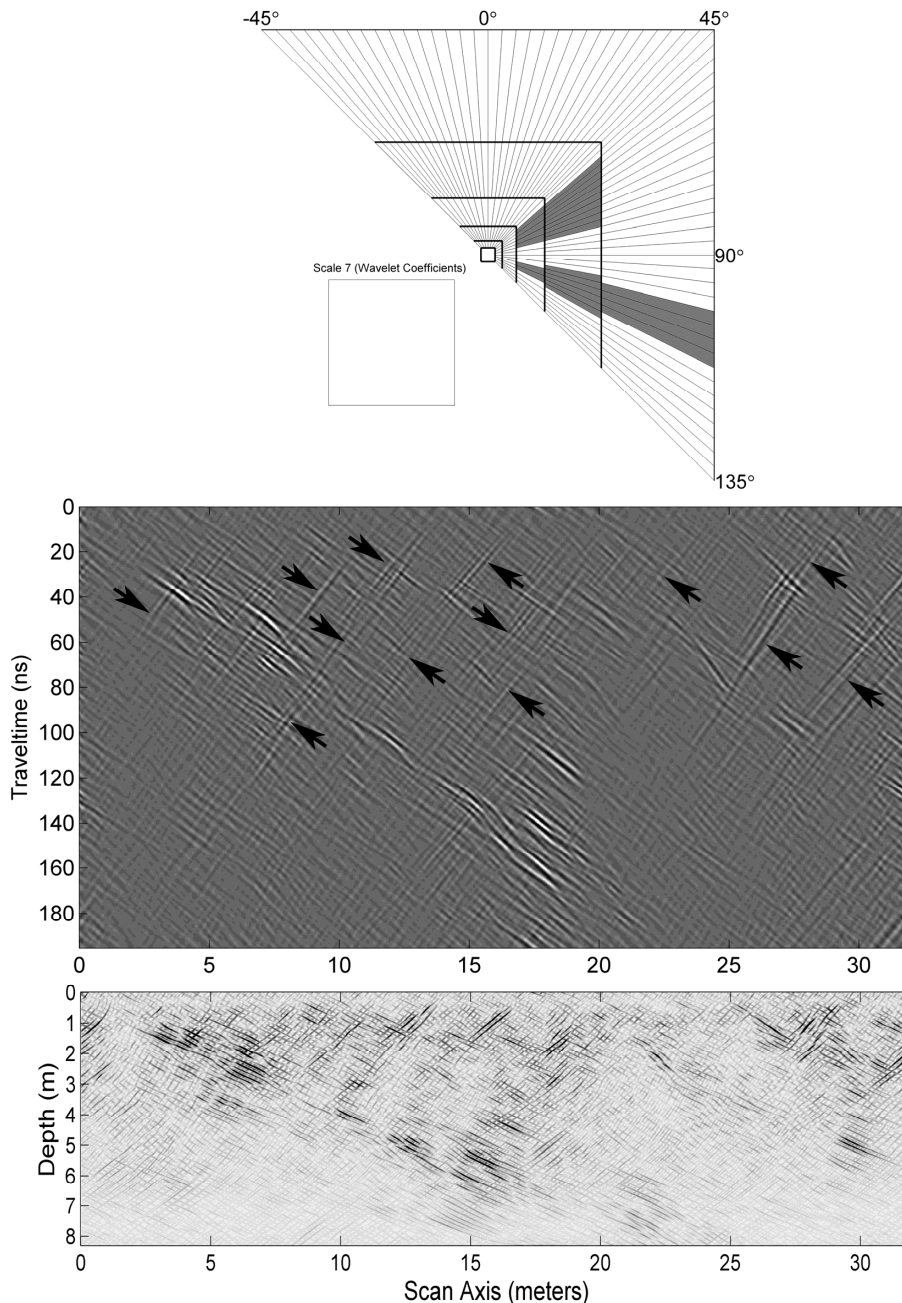


Figure 9d. The bottom panel illustrates a partial reconstruction of the radargram shown in Fig. 9b based of the subsets of coefficients $\{j = 4, l \in [27, 28]\} \cup \{j = 5, l \in [52, 57]\} \cup \{j = 6, l \in [53, 56]\}$ and $\{j = 4, l \in [18, 22]\} \cup \{j = 5, l \in [35, 44]\}$, as shown at the top. The former subset isolates reflections from small-aperture up-dipping antithetic fractures. The latter isolates reflections from down-dipping synthetic (main) fractures. The black arrows pointing to the up-dipping reflections are placed identically as in Fig. 9b (bottom).




Figure 10. (a) Example radargram obtained at the crest of Mt. Ktenias, Greece. The data has been pre-processed as detailed in the text and is used courtesy of Mr P. Sotiropoulos, Terra-Marine Ltd (<http://terra-marine.gr>). **(b)** Map of the centroid frequency in MHz, computed using the S-transform.

The fifth and final example is similar to the third, with the difference that the propagation medium is more complex. Details on the geological setting of the survey can be found in Tzanis (2013) and only a brief recount will be given herein. The data was collected with a Måla GPR system and 250MHz antenna at the crest of Mt Ktenias, Greece, as part of a geotechnical survey to investigate ground conditions at a site where high rising wind power generators were to be built. The ground comprises dipping, thin-plated limestone with intercalations of argillaceous material, intensely fragmented by faulting and multi-phase jointing. Karstification has nucleated at the loci of faults and major joints, with many cracks and karstic voids subsequently filled with lateritic material. Fig. 10a shows the example section after time-zero adjustment, global background removal, AGC with a 20ns Gaussian-tapered window, elimination of antenna self-clutter, resampling to a 1024 sample \times 1024 trace matrix (the sampling rate is 0.3818 ns and the trace spacing 0.0313 m) and, finally, f-k migration with a velocity of 0.0975 m/ns. The velocity was estimated by direct measurement in the field and the migration introduced widespread low amplitude/ low frequency/ short wavenumber spatio-temporal interference, more evident toward the latter times of the radargram.

The visual inspection of the data clearly indicates the existence of areas rich in low frequencies. The effect should be a consequence of the properties of the material filling these areas, i.e. preferential absorption of high frequencies in high attenuation domains. The analysis of the centroid frequency (Fig. 10b) indicates the existence of a complex propagation medium in which there are areas of severe high-frequency attenuation and areas of low, or altogether low attenuation, mainly observed below the 100ns line, from which there are no reflections and the centroid frequency is heavily influenced by the characteristics of random noise. High frequency attenuation is particularly evident in a rather broad region between 12–28 m, which extends to at least 140ns between 20–25m, as well as in an up-dipping area at the far right of the section, collocated with a fault. No gradual downshift of the centroid frequency was observed, implying that dispersion is not strong.

Fig. 10c illustrates a partial reconstruction based on the subset $\{j = 4, l \in [1, 28] - [21, 22]\}$, which comprises frequencies 233 MHz – 460 MHz, comparable to significantly higher than the nominal frequency of the antenna. Likewise, Fig. 10d illustrates a partial reconstruction based on the subset $\{j = 3, l \in [1, 28] - [20, 23]\}$, which comprises frequencies 113 MHz – 233 MHz that are lower to significantly lower than the nominal frequency of the antenna. In both cases, the effects of purely spatial disturbance have been eliminated.

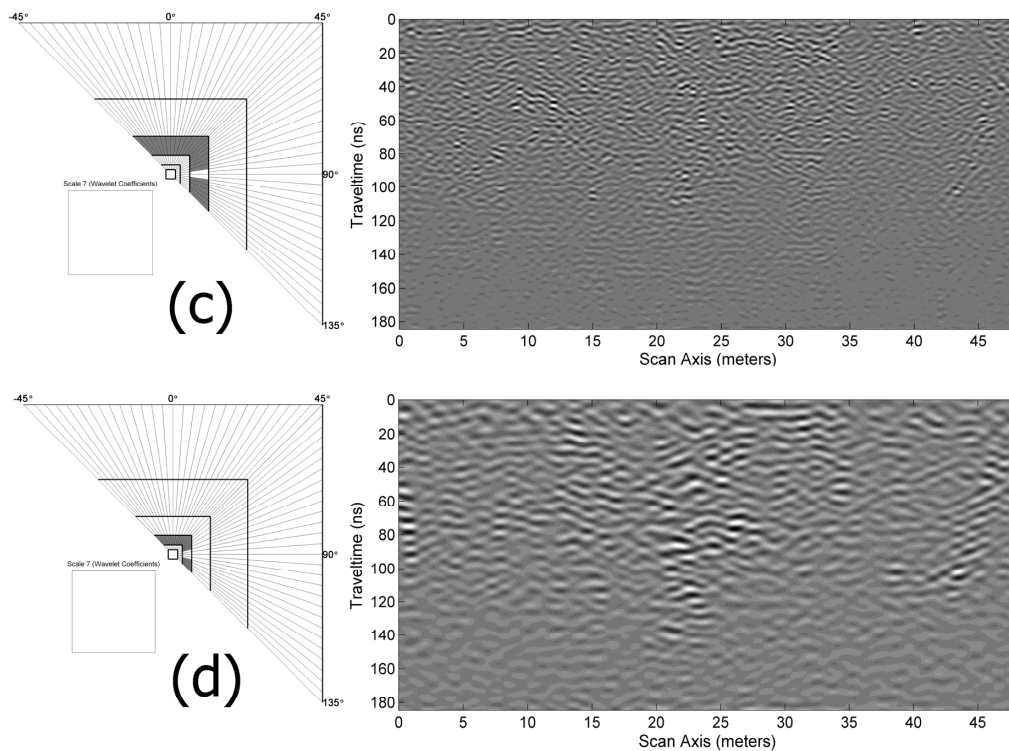


Figure 10. (c) The top-right panel illustrates a partial reconstruction of the radargram shown of Fig. 10a based on the subset of coefficients $\{j=4, l \in [1, 28] - [21, 22]\}$ shown in the top-left panel. (d) The bottom-right panel illustrates a partial reconstruction of the same radargram based on the subset $\{j = 3, l \in [1, 28] - [20, 23]\}$ shown in the bottom-left panel.

As expected, the results are very different and correspond to reflectors with very different material properties. Thus, Fig. 10c depicts sharper reflections originating at low attenuation interfaces: these are highly localized in time and rich in high frequencies so that they can be observed at short periods (finer temporal scales). Conversely, Fig. 10d depicts broadly localized dull reflections from high attenuation interfaces that absorb high frequencies and enrich the long period content of the signal. The dull reflectors

can be seen to comprise quasi-continuous anastomosing horizons, or to cluster in relatively narrow and vertically extended complexes. Given the geological situation of the study area, the former comprise interfaces between limestone fragments filled with high attenuation material, in which the larger and internally less damaged blocks appear in the form of lenses (hence the anastomoses). The latter (clusters) are signatures of cavities filled with high attenuation material. Examples of such structures can be observed between 20– 30 m and 20–150 ns, which is the signature of a karstic cavity beneath a bowl-shaped sinkhole, and the up-dipping fault at the far right of the section. In all cases, the dull reflectors are collocated with the low centroid frequency domains and the filling material is lateritic in composition. As mentioned above, at travel times longer than about 100ns the attenuation is low and the radargram is void of any significant reflection. Inasmuch as the signal penetrates to the corresponding depths –dull reflections from the karstic cavity are observed to at least 150ns – these areas comprise minimally damaged or even healthy rock, in which the bedding and tectonic interfaces have small apertures, are closed and/or dry and do not respond to the wavelengths of the 250MHz antenna. The fact that the filled cavity extends well into the healthy bedrock is evidence of its tectonic origin, i.e. that it nucleated in a fault or junction of faults.

4 Discussion and Conclusions.

The present work discusses a data processing method based on the Curvelet Transform (CT) and designed to facilitate the cancelation of complex noise and the retrieval of scale-and orientation dependent information from two-dimensional (B-scan) GPR data. The CT is a pyramidal multiscale and multidirectional decomposition (partition) of the f-k spectrum into highly anisotropic/ highly localized elements, expressed as an expansion that formulates an *optimally* sparse representation of the input data with *optimal* reconstruction property. This means that any recoverable piece of information with a given scale and orientation can be retrieved from a noisy data set (image), as if it was noise-free. By tuning the sampling rate, the order of the pyramidal decomposition (scale) and the order of the angular decomposition, one may resolve data components of at any scale and orientation, at least as fine as the order at which over- or under-sampling issues begin to hamper the computations (i.e. the limit of available information is exceeded).

The method is simple to use and very powerful. Herein it was shown that the CT can be used to:

- a) Enhance the S/N ratio by isolating and cancelling directional noise wavefronts of any scale and angle of emergence in such an adaptive way and precision, as not to interfere with recoverable data components.
- b) Isolate/ enhance and retrieve waveforms of specific scale and geometry for further scrutiny, as for instance distinguish signals from small and large aperture fractures and faults, different phases of fracturing and faulting, bedding etc.
- c) Exploit the temporal localization characteristics of the filter outputs to investigate the characteristics of signal propagation (hence material properties), albeit indirectly. This is possible because signal attenuation and temporal localization are closely associated. Thus, interfaces in low attenuation domains will tend to produce reflections richer in high frequencies and fine-scale localization, while interfaces in high attenuation domains will tend to produce dull reflections, rich in low frequencies and broad localization.

The method is also easy to automate; there are several tasks that can be routinely performed, as for instance suppression of the slow or very fast spatio-temporal variation associated with the coarser and

finest scales respectively and several types of (elaborate) f-k filtering. However, because the geometrical information contained in a data set is almost always unique, the real power of the CT rests on its exceptional capacity to adaptively manipulate it. Manual application is generally recommended, at least on a sample of representative radargrams: with a little experimentation by (educated) trial and error, the analyst may come up with optimal geometrical filtering schemes and, sometimes, operate with surgical precision. These schemes may then be applied routinely to other radargrams of the same data set.

The merits of the CT are apparent but, as with all methods, there are caveats. The case of incomplete partial reconstruction of transient wavefronts has been discussed in Section 3, where it was shown that low amplitude spurious oscillations can appear in cases of narrow-band operations on broadband events that are not characterized by a particular scale (e.g. transients). It is therefore necessary to exercise caution when treating data with such features. A possibly more serious issue is the risk of misinterpreting the output in terms of artefacts. For instance, when wavefronts of variable dip exist in the radargram and a highly anisotropic analysis is performed, e.g. targeting specific reflectors at specific orientations, the curvelets and the wavefronts will be matched only along the orientations specified by the curvelets. Depending on the lateral extent and angular content of the wavefronts, uniformly dipping localized features, or alignments of features may appear at different places of the output radargram. Notwithstanding that alignments often indicate the existence of intermittent reflective surfaces and that the filter performs as expected, it is easy to see how a false impression of artefacts can be created. In most cases, the false artefacts are actually subtle data components that are not immediately apparent to the observer. Caution, again is needed so as not to dismiss a valid output.

A second point to consider is that in GPR practice, the length of the radargram is frequently very long, e.g. hundreds of meters, so that the size of the spatial dimension (scan line) is very large when compared to the size of the temporal dimension. In such cases different subsurface reflectors with different geometrical characteristics (angular content) may appear at different places along the scan line. The latter is not an issue because the CT is optimally sparse representer and only a small set of scales and curvelets are usually needed in the reconstruction of a given geometrical object. Therefore, different reflections with different dips can always be reconstructed with the appropriate choice of scales and curvelets. However, it may not be practical to treat a very long radargram as a single image because the angular decomposition may be very awkward. Depending on the data at hand, an obvious way around this problem is to split the radargram in consecutive subsections and treat each subsection separately.

A final valid and important question is of whether other methods exist, that can offer analogous or even improved efficiency. The answer is not simple and also not negative, as will be discussed below. To the best of the Author's knowledge, there are only two other methods designed to retrieve geometrical information from GPR (and seismic reflection) data. The first is a parametric local-dip image decomposition approach that utilizes Local Radon Transforms (Theune et al, 2006a; Theune et al, 2006b). Although applications of the method are not common in the literature, it has been shown to be robust and successful in extracting coherent dipping features. It can also be used to remove coherent and random noise, separate wavefields and reconstruct missing data, although its efficiency to this effect has not been evaluated. In addition, the method cannot separate events of different scale and the range of dips it can handle is limited by the (inherent) risk of aliasing in the tau- p transform.

Another simple and intuitive filtering scheme is the (multi)directional filter approach described in Tzanis (2013). This is different in design and operation and there is no objective metric by which to compare it quantitatively to the CT. The relative performance of the two methods can only be assessed qualitatively, in terms of the similitude of their outputs for the same input data and where the degree of similitude is a matter of common sense. The data of the first and second examples presented in Section 4 have also been

treated by Tzanis (2013) and can be a basis for comparisons with simple data sets. Analogous direct comparisons for complex data sets can be found in Tzanis (2014).

Depending on the data and the objective of the analysis, either method has properties that may favour its application over the other. In the CT for instance, localization in *both* frequency and wavenumber is somewhat limited by the order of the pyramidal decomposition. Conversely, it has very versatile angular resolution characteristics and may resolve features with small differences in their angular distribution. In the directional filtering scheme, localization and angular resolution depend on the type and scale (length) of the wavelet basis and the anisotropy of the filter, so that it can vary from broad to very fine. Multidirectional filtering is applied by combining the outputs of the same directional filter rotated to different angles under adaptive control, so that it remains *tuned* at a given frequency or wavenumber. In this way the filter can be localized on specific frequency or spatial aspects of the data, a trait that may not always be possible with a pyramidal decomposition.

It is rather easy to verify that for simple linear or quasi-linear reflectors, such as the one shown in the first example of Section 4, the two methods compare very well, particularly when the anisotropy of the directional filter is similar to the anisotropy of the curvelets. If wavelets with intermediate localization properties are used for multidirectional filtering, the results are generally very comparable for data of moderate complexity. However, the possibility to fine tune the multidirectional filters and to choose the extent of frequency or spatial localization may be advantageous in resolving geometric information in very complex data. Conversely, high localization is not recommended for data with broadband geometric features, while the versatile angular decomposition of the CT facilitates the discrimination of wavefronts with small differences in their orientation. Moreover, the elaborate filtering schemes afforded by the CT enable the successful processing of data swamped in complex noise. For example, it should not take long to verify that the data of Fig. 8a could not have been treated effectively with the simple spectral structures of the (multi)directional filters!

The CT is, in effect, a very elaborate f-k filtering technique. In some cases with simple data and/or noise structures, conventional f-k filters may perform equally well. However, when the structural and/or noise complexity increases, conventional f-k filtering techniques may be very impractical. The advanced orientation-sensitive X-let transforms may also offer efficient alternatives, depending of course, on their design. Some older designs are certainly not as efficient. Ridgelets, for instance, were intended to represent objects with discontinuities on straight lines. Other X-lets optimized for the processing of edges exist, such as shearlets, riplelets and contourlets; most of them are genetically related to the CT, have similar approximation rates and may perform analogously well. When the data is strongly oscillatory, designs optimized to process oscillating textures (e.g. wave atoms) may be superior because “curvelets only capture the coherence of the pattern along the oscillations, not across” (Demanet and Ying, 2007). Adaptive low-dimensional approximations (e.g. wedgelets) and adaptive orthogonal expansions (e.g. bandlets) may perform very efficient approximations of objects with geometric singularities; however, it is doubtful whether such objects are common in GPR data. The list is certainly long and cannot be exhausted herein, to anyone’s satisfaction. Some of the advanced X-lets, specifically of the type optimized to process oscillatory patterns, will be examined in a follow-up research.

In conclusion, the CT offers a provably effective and powerful method to process and analyze noisy and complex data, with some limitations arising when the data is strongly oscillatory, or if very specific traits (in frequency or wavenumber) need to be resolved. These limitations can be addressed with other methods, as for instance with multidirectional filtering techniques as detailed above. In the end, the CT and multidirectional filters both provide a powerful and complementary set of tools with which to retrieve almost any recoverable structural trait from GPR data.

REFERENCES

- Bano, M., 1996. Constant dielectric losses of ground-penetrating radar waves, *Geophys. J. Int.*, **124**, 279-288.
- Bano, M., 2006. Effects of the transition zone above a water table on the reflection of GPR waves, *Geophys. Res. Lett.*, **33**, L13309; doi:10.1029/2006GL026158.
- Bano, M., 2007. Modelling GPR reflections from the transition zone above the water table, analytical approach in frequency domain, SEG Technical Program Expanded Abstracts 2007: pp. 1192-1196; doi: 10.1190/1.2792719.
- Cagnoli, B. and Ulrych, T.J., 2001. Singular value decomposition and wavy reflections in ground-penetrating radar images of base surge deposits, *Journal of Applied Geophysics*, **48**, 175 – 182.
- Candès, E., 1999. Harmonic analysis of neural networks. *Appl. Comput. Harmon. Anal.*, **6**, 197-218.
- Candès, E. and Donoho, D., 1999. Ridgelets: A key to higher-dimensional intermittency? *R. Soc. Lond. Philos. Trans. Ser. A Math. Phys. Eng. Sci.*, **357**, 2495-2509.
- Candès, E. and Donoho, D., 2003a. Continuous curvelet transform: I. Resolution of the wavefront set. *Appl. Comput. Harmon. Anal.*, **19**, 162-197.
- Candès, E. and Donoho, D., 2003b. Continuous curvelet transform: II. Discretization and frames. *Appl. Comput. Harmon. Anal.*, **19**, 198-222.
- Candès, E. and Donoho, D., 2004a. New tight frames of curvelets and optimal representations of objects with piecewise C^2 singularities. *Comm. Pure Appl. Math.*, **57**, 219-266.
- Candès, E. and Donoho, D., 2004b. Curvelets: new tools for limited angle tomography, Technical Report, California Institute of Technology.
- Candès, E. J., and Demanet L., 2005. The curvelet representation of wave propagators is optimally sparse, *Comm. Pure Appl. Math.*, **58** (11), 1472–1528.
- Candès, E. J., L. Demanet, D. L. Donoho, and L. Ying, 2006. Fast discrete curvelet transforms (FDCT). *Multiscale Modeling and Simulation*, **5**, 861–899.
- Chui, C.K., 1992. *An introduction to wavelets*, Academic Press, New York.
- Deighan, A. J., and Watts, D. R., 1997. Ground-roll suppression using the wavelet transform. *Geophysics*, **62**, 1896–1903.
- Demanet, L. and Ying, L., 2007. Wave Atoms and Sparsity of Oscillatory Patterns. *Appl. Comput. Harmon. Anal.*, **23** (3), 368-387.
- Do, M. and Vetterli, M., 2005. The contourlet transform: an efficient directional multiresolution image representation. *IEEE Trans. Image Process.*, **14** (12), 2091-2106.
- Donoho, D. and Huo, X., 2002. Beamlets and multiscale image analysis, in: T. Barth et al. (Eds.), *Multiscale and Multiresolution Methods*, Springer Lecture Notes in Comput. Sci. Eng., vol. **20**, pp. 149-196.
- Donoho, D., 1999. Wedgelets: nearly minimax estimation of edges. *Ann. Statistics*, **27** (3), 859-897.
- Feichtinger, H.G. and Strohmer, T., 1998. *Gabor Analysis and Algorithms*, Birkhäuser; ISBN 0817639594
- Feichtinger, H.G. and Strohmer, T., 2003. *Advances in Gabor Analysis*, Birkhäuser; ISBN 0817642390.
- Freeman, W. and Adelson, E., 1991. The design and use of steerable filters. *IEEE Trans. Pattern Analysis and Machine Intelligence*, **13** (9), 891-906.
- Fukunaga, K., 1990. *Introduction to Statistical Pattern Recognition*, 2nd Ed., Academic Press, London.

- Grigorescu, C., Petkov, N. and Westenberg, M.A., 2003. Contour detection based on nonclassical receptive field inhibition. *IEEE Trans. Image Process.*, **12** (7), 729-739.
- Guo, K. and Labate, D., 2007. Optimally sparse multidimensional representation using shearlets, *SIAM J. Math. Anal.*, **39**, 298–318.
- Irving, J.D. and Knight, R.J., 2003. Removal of wavelet dispersion from ground-penetrating radar data. *Geophysics*, **68** (3), 960-970.
- Irving, J. and Knight, R., 2006. Numerical modeling of ground-penetrating radar in 2-D using MATLAB, *Computers & Geosciences*, **32**, 1247–1258; doi: 10.1016/j.cageo.2005.11.006
- Jacques, L., Duval, L., Chaux, C. and Peyré, G., 2011. A panorama on multiscale geometric representations, intertwining spatial, directional and frequency selectivity, *Signal Processing*, **91**, 2699-2730; doi:10.1016/j.sigpro.2011.04.025
- Jeng, Yih, Lin, C-H., Li, Y-W., Chen, C-S. and Huang, H-H., 2009. Application of multiresolution analysis in removing ground-penetrating radar noise. *Frontiers+Innovation – 2009 CSPG-CSEG-CWLS Convention*, 416-419.
- Lee, T., 2008. Image representation using 2D Gabor wavelets. *IEEE Trans. Pattern Anal.*, **18** (10), 1-13.
- Leblanc, G. E., Morris, W. A., and Robinson, B., 1998. Wavelet analysis approach to de-noising of magnetic data. *SEG Expanded Abstract*, 554-557.
- Lu, Y. and Do, M., 2007. Multidimensional directional filter banks and surfacelets. *IEEE Trans. Image Process.*, **16** (4), 918-931.
- Lucius, J.E. and Powers, M.H., 2002. GPR Data Processing Computer Software for the PC. USGS Open-File Report 02-166.
- Ma, J. and Plonka, G., 2010. The Curvelet Transform, *IEEE Signal Processing Magazine*, **118**; doi: 10.1109/MSP.2009.935453
- Mallat, S. and Peyré, G., 2007. A review of bandlet methods for geometrical image representation, *Numer. Algorithms*, **44** (3), 205-234.
- Mallat, S.G., 1999. *A Wavelet Tour of Signal Processing*, Academic Press.
- Matos, M. D. and Osorio, P. M., 2002. Wavelet transform filtering in the 1D and 2D for ground roll suppression. *SEG Expanded Abstract*, 2245-2248.
- Miao, X., and Cheadle, S. P., 1998. Noise attenuation with wavelet transforms. *SEG Expanded Abstract*, 1072-1075.
- Nuzzo, L., and Quarta, T., 2004. Improvement in GPR coherent noise attenuation using $\tau - p$ and wavelet transforms. *Geophysics* **69**, 789–802.
- Roberts, A. (2001) Curvature attributes and their application to 3D interpreted horizons. *First Break*, **19** (2), 85–100; doi: 10.1046/j.0263-5046.2001.00142.x.
- Rudzki, M., 2008, Coherent Noise Attenuation in the GPR Data via the Karhunen-Loève Transform, Proceedings, 14th European meeting of the EAGE, Krakow, Poland, 29-30 September 2008.
- Simoncelli, E., Freeman, W., Adelson, E. and Heeger, D., 1992. Shiftable multiscale transforms. *IEEE Trans. Inform. Theory*, **38** (2), 587-607.
- Smith, H. F., 1998. A Hardy space for Fourier integral operators, *J. Geom. Analysis*, **7**, 629 – 653.
- Stockwell, R.G., Mansinha, L. and Lowe, R.P., 1996. Localization of the complex spectrum: The S-transform. *IEEE Trans. Signal Proces.*, **44**, 998-1001.

- Sun, B., Ma, J., Chauris, H. and Yang, H., 2010. Solving the Wave Equation Using Curvelets, Proceedings, 72nd EAGE Conference & Exhibition, Barcelona, Spain June 14 - 17, 2010
- Stoffa, P.L., Fokkema, J.T., de Luna Freire, R.M. and Kessinger, W.P., 1990. Split-step Fourier migration. *Geophysics*, **55**, 410-421.
- Theune, U., Sacchi, M.D. and Schmitt, D.R., 2006a. Least-squares local Radon Transforms for dip-dependent GPR image decomposition, *Journal of Applied Geophysics*, **59**, 224-235; doi: 10.1016/j.jappgeo.2005.10.003.
- Theune, U., Rokosh, D., Sacchi, M.D. and Schmitt, D.R., 2006b. Mapping fractures with GPR: a case study from Turtle Mountain. *Geophysics*, **71** (5), B139-B150.
- Tzanis, A., 2010. matGPR Release 2: A freeware MATLAB[®] package for the analysis and interpretation of common and single offset GPR data, *FastTimes*, **15** (1), 17 – 43.
- Tzanis, A., 2013. Detection and extraction of orientation-and-scale-dependent information from two-dimensional GPR data with tuneable directional wavelet filters. *Journal of Applied Geophysics*, **89**, 48-67; doi: 10.1016/j.jappgeo.2012.11.007.
- Tzanis, A., 2014. Signal enhancement and geometric information retrieval from 2-D GPR data with multiscale, orientation-sensitive filtering methods, *First Break*, **32** (8), 91-98.
- Xie, X., Zeng, C. and Wang, Z., 2013, GPR signal enhancement using band-pass and K–L filtering: a case study for the evaluation of grout in a shielded tunnel, *J. Geophys. Eng.* 10 034003.
- Xu, J., Yang, L. and Wu, D., 2010. Ripplet: a new transform for image processing, *J. Vis. Commun. Image Repr.*, **21** (7), 627–639
- Zhao, A., Jiang, Y. and Wang, W., 2005. Signal-to-noise ratio enhancement in multichannel GPR data via the Karhunen-Loeve Transform, *Proceedings, Progress in Electromagnetic Research 2005*, Hangzhou, China, August 22-26 2005.