

Wavelet-based Segmentation on the Sphere

Xiaohao Cai*, Christopher G. R. Wallis*, Jennifer Y. H. Chan*, and Jason D. McEwen*

*Mullard Space Science Laboratory (MSSL), University College London (UCL), UK
x.cai@ucl.ac.uk, chris.wallis@ucl.ac.uk, y.chan.12@ucl.ac.uk, jason.mcewen@ucl.ac.uk

Abstract

Segmentation, a useful/powerful technique in pattern recognition, is the process of identifying object outlines within images. There are a number of efficient algorithms for segmentation in Euclidean space that depend on the variational approach and partial differential equation modelling. Wavelets have been used successfully in various problems in image processing, including segmentation, inpainting, noise removal, super-resolution image restoration, and many others. Wavelets on the sphere have been developed to solve such problems for data defined on the sphere, which arise in numerous fields such as cosmology and geophysics. In this work, we propose a wavelet-based method to segment images on the sphere, accounting for the underlying geometry of spherical data. Our method is a direct extension of the tight-frame based segmentation method used to automatically identify tube-like structures such as blood vessels in medical imaging. It is compatible with any arbitrary type of wavelet frame defined on the sphere, such as axisymmetric wavelets, directional wavelets, curvelets, and hybrid wavelet constructions. Such an approach allows the desirable properties of wavelets to be naturally inherited in the segmentation process. In particular, directional wavelets and curvelets, which were designed to efficiently capture directional signal content, provide additional advantages in segmenting images containing prominent directional and curvilinear features. We present several numerical experiments, applying our wavelet-based segmentation method, as well as the common K-means method, on real-world spherical images, including an Earth topographic map, a light probe image, solar data-sets, and spherical retina images. These experiments demonstrate the superiority of our method and show that it is capable of segmenting different kinds of spherical images, including those with prominent directional features. Moreover, our algorithm is efficient with convergence usually within a few iterations.

Keywords:

image segmentation, wavelets, curvelets, tight frame, sphere

1. Introduction

Spherical images are common in nature, for example, in cosmology [37], astrophysics [50], planetary science [3], geophysics [52], and neuro-science [45], where images are naturally defined on the sphere. Clearly, images defined on the sphere are different to Euclidean images in 2D and 3D in terms of symmetries, coordinate systems and metrics constructed (see for example [29]). Image segmentation aims to separate a given image into different components, where each part shares similar characteristics in terms of, e.g., edges, intensities, colours, and textures. It generally serves as a preliminary step for object recognition and interpretation, and is a fundamental yet challenging task in image processing. In this paper, we present an effective segmentation method that uses spherical wavelets to segment spherical images.

In the literature, many different approaches have been proposed for image segmentation for 2D, 3D and vector-valued images, e.g., [10, 11, 30, 40, 46, 51, 53]. In particular, in [40] the well-known Mumford-Shah model was proposed, which formulates the image segmentation problem by minimising an energy function and finding optimal piecewise smooth approximations of the given image. More details about these kind of methods can be found in [6, 8, 13, 17]. These types of methods generally give good segmentation results. However, their applicabilities and performance heavily depend on the models used; in some cases (e.g. segmenting images containing complex textures) the models are difficult or expensive to compute due to the non-convex nature of the problem. In [51], a graph-cut based method was proposed to segment point clouds into different groups. The more pixels in the image, the larger the

size of the eigenvalue problem needs to be solved, which makes the method inefficient in terms of speed and accuracy. Methods based on deformable models [23, 30] segment via evolving geodesic active contours that are built from a partial differential equation, with the ability to detect twisted, convoluted and occluded structures, but are sensitive to noises and blurs in images.

Recently, segmentation methods [8, 13, 14] designed utilising techniques in image restoration (e.g. [43, 49]), were proposed for gray-scale images. In [8], a segmentation model that combines the image segmentation model of [40] and the data fidelity terms from image restoration models [43, 49] was considered to deal with images contaminated by different types of noise (e.g. Gaussian, Poisson or impulsive noise). In [13], the methodology of two-stage methods, solving image restoration models first followed by a thresholding second stage, was proposed. These methods were later named SaT (smoothing and thresholding) segmentation methods. One advantage of the SaT methods is the fast speed of implementation. Akin to the SaT methods, the T-ROF method (thresholding the Rudin-Osher-Fatemi model) in [8] concluded that the thresholding approach for segmentation was equivalent to solving the Chan-Vese segmentation model [17]; more detailed theoretical proofs can be found in [12]. Based on the SaT methodology, a method named SLaT [11] was proposed for degraded colour image segmentation.

In addition to the aforementioned methods, approaches based on wavelets and tight frames [4, 9, 10] have been proposed for segmentation. In [9, 10], a tight-frame based segmentation method was designed for a vessel segmentation problem in medical imaging. The major advantage of this method is the ability to segment twisted, convoluted and occluded structures without user interactions. Moreover, the ability of the method to follow the branching of different layers, from thinner to larger structures, makes the method a good candidate for a tubular-structured segmentation problem in medical imaging. However, all the tight-frame systems discussed and used in [10] (e.g. framelets [47], contourlets [21], curvelets [15, 16], and dual-tree complex wavelet¹) are designed for 2D or 3D data on a Euclidean manifold. Consequently, these approaches cannot be applied to problems where data-sets live natively on the sphere.

Wavelets have become a powerful analysis tool for spherical images, due to their ability to simultaneously extract both spectral and spatial information. A variety of wavelet frameworks have been constructed on the sphere in recent years, e.g. [2, 5, 32, 34, 36], and have led to many insightful scientific studies in the fields mentioned above (see [37, 50, 3, 52, 45]). Different types of wavelets on the sphere have been designed to probe different structure in spherical images, for example isotropic or directional and geometrical features, such as linear or curvilinear structures, to mention a few. Axisymmetric wavelets [42, 5, 28, 54] are useful for probing spherical images with isotropic structure, directional wavelets [36, 56, 32, 34] for probing directional structure, ridgelets [39, 54] for analysing antipodal signals on the sphere, and curvelets [54, 18] for studying highly anisotropic image content such as curve-like features (we refer to [15, 16] for the general definition of Euclidean ridgelets and curvelets). Fast algorithms have been developed to compute exact forward and inverse wavelet transforms on the sphere for very large spherical images containing millions of pixels [36, 33, 34] (leveraging novel sampling theorems on the sphere [38] and the rotation group [31]). Localisation properties of wavelet constructions have also been studied in detail [5, 42, 32], showing important quasi-exponential localisation and asymptotic uncorrelation properties for certain wavelet constructions. An investigation into the use of axisymmetric and directional wavelets for sparse image reconstruction was performed recently in [55], showing excellent performance. The spherical wavelets adopted in the experimental section of this paper are reviewed briefly in Section 2.

In this paper, we devise an iterative framework for segmenting spherical images using wavelets defined on the sphere, extending the method proposed in [9, 10]. The first stage of the method, as a preprocessing step, suppresses noises in the given data by soft thresholding wavelet coefficients. Then, potential boundary pixels are classified gradually by the iterative procedure. The framework is compatible with any arbitrary type of spherical wavelet, such as the axisymmetric wavelets, directional wavelets, or curvelets mentioned above. The iterative strategy in the proposed framework is effective, particularly for images containing anisotropic textures. There is also flexibility regarding the implementation of iterations. Motivated by the SaT methodology in [14, 13], when segmenting images containing many (or mostly) isotropic structures, the iterative strategy in our method can be replaced by a simple thresholding to reduce the computation time significantly without sacrificing segmentation quality considerably. We test the proposed framework on a variety of types of spherical images, including an Earth topographic map, a light probe (spherical) image, two sets of solar data, and two retina images projected on the sphere.

¹<http://taco.poly.edu/WaveletSoftware/>

To the best of our knowledge, this is the first segmentation method that works directly on the *whole* sphere and is practical for any type of spherical images, benefiting from the compatibility of the method with any type of spherical wavelets. A method [48] was proposed for segmenting spherical particles in volumetric data sets based on an extension of the generalised Hough transform and an active contour approach. However, the data considered in [48] were 3D data containing spherical-like particles, not data defined on the sphere directly. Another method [41] was proposed for 3D shape segmentation based on an active contour formulation with a shape prior which was formed by using spherical wavelets. However, the data considered in [41] were 3D data containing 3D shapes, not data that are defined natively on the sphere.

Very recently, spherical deep neural networks, an important tool in artificial intelligence, have emerged, e.g., SphereNet [20], Spherical CNNs [19, 22], Spherical CNNs on Unstructured Grids [26], and DeepSphere [44]. The success of deep learning generally depends on the exploitation of data properties by the network architectures for efficient and principled learning, as well as on the quality of the training sets. Our proposed method tackles the task of segmentation with a completely different approach and does not require any training sets. On the other hand, an extension of our method with deep neural networks is possible (see a brief discussion in Section 3).

The main contributions in this paper are: (1) a segmentation framework for spherical images is devised, for the first time; (2) the framework uses an iterative strategy with the flexibility to tailor the iterative procedure according to data types and features; (3) spherical wavelets, including axisymmetric wavelets, directional wavelets and the newly-constructed hybrid directional-curvelet wavelets, are implemented and tested in the framework; (4) a series of applications are presented, illustrating the performance of our proposed segmentation method.

The remainder of this paper is organised as follows. In Section 2, we review related work about spherical wavelets and segmentation methods, and present our new hybrid directional-curvelet wavelet construction. In Section 3, we introduce our spherical segmentation method. In Section 4, the proposed method and methods for comparison are tested on a variety of spherical images such as an Earth map, light probe images, and two solar maps. To further demonstrate the ability of our method on segmenting highly directional and elongated structures, in Section 4 we also apply it to retinal images, which contain a complex network of blood vessels. Conclusions are given in Section 5.

2. Background

Let $f \in L^2(\mathbb{S}^2)$ be the given image defined on the sphere \mathbb{S}^2 . Without loss of generality, we assume f in $[0, 1]$. Let $\omega = (\theta, \phi) \in \mathbb{S}^2$ denote spherical coordinates with colatitude $\theta \in [0, \pi]$ and longitude $\phi \in [0, 2\pi)$. Let \mathbb{S}^2 be the discretised sphere of \mathbb{S}^2 . We review sampling, wavelets and discrete gradient operators on the sphere subsequently, before recalling the tight-frame based segmentation method of [9, 10]. In addition, we present a new hybrid directional-curvelet wavelet construction.

2.1. Sampling on the sphere

We adopt the equiangular sampling theorem on the sphere of [38], which defines how to capture the information content in a signal band-limited at L in $\sim 2L^2$ samples. This sampling theorem requires the fewest number of samples to capture all information content on band-limited spherical images. In addition, fast algorithms to perform the associated spherical harmonic transform are presented [38]. Typically, we consider band-limited spherical images whose spherical harmonic coefficients $f_{\ell m} = 0, \forall \ell \geq L$, where $f_{\ell m} = \langle f, Y_{\ell m} \rangle$ and $Y_{\ell m} \in L^2(\mathbb{S}^2)$ are the spherical harmonics, with $\ell \in \mathbb{N}$ and $m \in \mathbb{Z}$ satisfies $|m| \leq \ell$. In practice many real-world signals can be approximated accurately by a band-limited signal). The equiangular sample positions of the sphere associated with this sampling theorem are given by

$$\theta_t = \frac{\pi(2t+1)}{2L-1}, \quad \phi_p = \frac{2\pi p}{2L-1}$$

where $t \in \{0, 1, \dots, L-1\}$ and $p \in \{0, 1, \dots, 2L-2\}$ index the equiangular samples in θ and ϕ , respectively. For example, when $L = 512$, the sphere is discretised with $512 \times 1023 = 523776$ samples. Please refer to [38] and references therein for more information about sampling on the sphere.

2.2. Wavelets on the sphere

In many real-life problems, data to be processed are usually in a discretised form, as described previously. Also, exact reconstruction of the signal is commonly desired. Scale-discretised wavelets [18, 28, 34, 32, 36, 56] on the sphere allow the exact synthesis of discrete spherical images from their wavelet coefficients. We adopt scale-discretised wavelet constructions in this work, which we review concisely in this section. In addition, we present a new hybrid directional-curvelet scale-discretised wavelet construction.

Wavelet transforms. Let $\Psi^{(j)} \in L^2(\mathbb{S}^2)$ be the wavelet with wavelet scales $j \in \mathbb{N}$ and $0 \leq J_{\min} \leq j \leq J_{\max}$, which encode the angular localisation of $\Psi^{(j)}$, where J_{\min} and J_{\max} are the minimum and maximum wavelet scales considered, respectively; see [34] for more details about j . For directional wavelet transforms, wavelet coefficients are defined on the rotation group $\text{SO}(3)$, parameterised by Euler angles $\rho = (\alpha, \beta, \gamma) \in \text{SO}(3)$ with $\alpha \in [0, 2\pi)$, $\beta \in [0, \pi]$ and $\gamma \in [0, 2\pi)$. Wavelet coefficients $W^{\Psi^{(j)}} \in L^2(\text{SO}(3))$ are computed by the wavelet forward transform (analysis) defined by

$$W^{\Psi^{(j)}}(\rho) \equiv (f \circledast \Psi^{(j)})(\rho) \equiv \langle f, \mathcal{R}_\rho \Psi^{(j)} \rangle = \int_{\mathbb{S}^2} d\Omega(\omega) f(\omega) (\mathcal{R}_\rho \Psi^{(j)})^*(\omega), \quad (1)$$

where \mathcal{R}_ρ is a rotation operator related to a 3D rotation matrix R_ρ by $(\mathcal{R}_\rho \Psi^{(j)})(\omega) \equiv \Psi^{(j)}(R_\rho^{-1} \hat{\omega})$ ($\hat{\omega}$ is the Cartesian vector of ω), $d\Omega(\omega) = \sin\theta d\theta d\phi$ is the usual rotation invariant measure on the sphere; the symbol $\langle \cdot, \cdot \rangle$, the operator \circledast , and $*$ denote the inner product of functions, directional convolution on the sphere and complex conjugation, respectively. Low-frequency content of the signal not probed by wavelets are probed by the scaling function $\Phi \in L^2(\mathbb{S}^2)$, which is generally axisymmetric. The scaling coefficients $W^\Phi \in L^2(\mathbb{S}^2)$ are given by

$$W^\Phi(\omega) \equiv (f \odot \Phi)(\omega) \equiv \langle f, \mathcal{R}_\omega \Phi \rangle = \int_{\mathbb{S}^2} d\Omega(\omega') f(\omega') (\mathcal{R}_\omega \Phi)^*(\omega'), \quad (2)$$

where $\mathcal{R}_\omega = \mathcal{R}_{(\phi, \theta, 0)}$, and the operator \odot denotes axisymmetric convolution on the sphere.

The spherical image f can be synthesised perfectly from its wavelet and scaling coefficients (under the wavelet admissibility condition [34]) by the wavelet backward transform (synthesis) by

$$f(\omega) = \int_{\mathbb{S}^2} d\Omega(\omega') W^\Phi(\omega') (\mathcal{R}_{\omega'} \Phi)(\omega) + \sum_{j=J_{\min}}^{J_{\max}} \int_{\text{SO}(3)} d\varrho(\rho) W^{\Psi^{(j)}}(\rho) (\mathcal{R}_\rho \Psi^{(j)})(\omega), \quad (3)$$

where $d\varrho(\rho) = \sin\beta d\alpha d\beta d\gamma$ is the usual invariant measure on $\text{SO}(3)$.

Construction of different types of wavelets. Spherical wavelets, constructed to ensure the admissibility condition is satisfied, are defined in harmonic space in the factorised form by

$$\Psi_{\ell m}^{(j)} \equiv \sqrt{\frac{2\ell+1}{8\pi^2}} \kappa^{(j)}(\ell) \zeta_{\ell m}, \quad (4)$$

where kernel $\kappa^{(j)} \in L^2(\mathbb{R}^+)$, a positive real function, is constructed to be a smooth function with compact support to control the angular localisation properties of wavelet $\Psi^{(j)}$, with harmonic coefficients $\Psi_{\ell m}^{(j)} = \langle \Psi^{(j)}, Y_{\ell m} \rangle$; see [32] for the detailed definition. The directionality component $\zeta \in L^2(\mathbb{S}^2)$, with harmonic coefficients $\zeta_{\ell m} = \langle \zeta, Y_{\ell m} \rangle$, is designed to control the directional localisation properties of $\Psi^{(j)}$. The wavelets recovered are steerable when imposing an azimuthal band-limit N on the directionality component such that $\zeta_{\ell m} = 0$ for $|m| \geq N, \forall \ell, m$. While steerability is achieved, the directional localisation of the wavelet is controlled by imposing a specific form for the directional auto-correlation of the wavelet. The detailed construction of ζ and $\zeta_{\ell m}$ for directional wavelets can be found in [32] and those for curvelets can be found in [18]. In particular, the spherical curvelets proposed in [18] exhibits the parabolic scaling relation. Such a geometric feature is unique to curvelets, making it highly anisotropic and directionally sensitive, and thus suitable for extracting local curvilinear structures effectively. Moreover, scale-discretised wavelets support the exact analysis and synthesis of both scalar and spin signals, although only the former are considered herein.

Fig. 1 and Fig. 2 show the harmonic tilings of different types of scale-discretised wavelets and the corresponding wavelets plotted on the sphere, respectively. We refer the reader to [28, 32, 36, 34, 56], and [18] for details about the

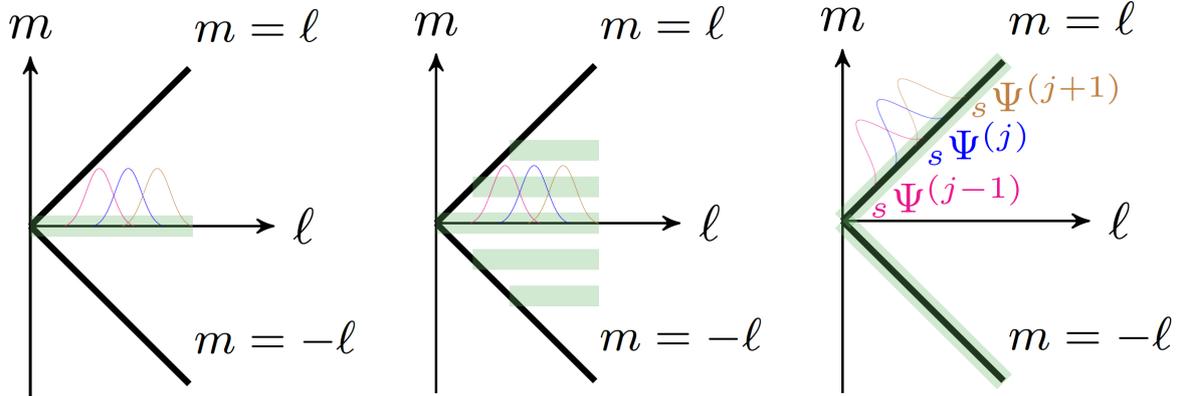


Figure 1: Harmonic tilings of different types of wavelets, including axisymmetric wavelets, directional wavelets, and curvelets, respectively, from left to right (refer to [18]).

construction of scale-discretised axisymmetric and directional wavelets, and curvelets, respectively. Code to compute these wavelet transforms is public and available in the existing S2LET² package, which relies on the SSHT³ [38] and SO3⁴ [31] packages.

Hybrid wavelets. Different wavelet transforms have differing computational requirements. Generally, computing axisymmetric wavelet transforms are fastest with computational time scaling as $O(L^3)$ [28], directional wavelet transforms are slower with computational time scaling as $O(NL^3)$ [34], while curvelet transforms are the slowest with computational time scaling as $O(L^3 \log_2 L)$ [18]. As an example, a complete round-trip of a forward and backward wavelet transform, with band-limit $L = 512$, takes a few seconds on a Macbook with i5 processor for axisymmetric wavelets, several minutes for directional wavelets and several hours for curvelets (see Tables 1, 2, 3 and 4). However, the increase in computational cost is offset by an improved ability to represent directional and curvilinear structure for directional wavelets and curvelets, respectively.

With the aim to exploit the advantages of the curvelet transform [18] while shortening the computational time needed, here we construct a hybrid form of wavelet transform on the sphere using both curvelets and directional wavelets. The idea (proposed as a future work in our paper [18]) is to describe small-scale features with directional wavelets and remaining features with curvelets, thereby inheriting the excellent directional localisation of curvelets and computational advantages of directional wavelets.

The separation between the two wavelet types is performed in harmonic space, at a defined transition band-limit L_{trans} . The curvelet transform is performed up to the band-limit L_{trans} , ignoring the final wavelet scale. This provides the large-scale curvelet coefficients. In order to calculate the component of the image represented by curvelets, the inverse transform is performed, yielding f^{curv} . The directional wavelet coefficients are found by first subtracting this image from the original, $f^{\text{dir}} = f - f^{\text{curv}}$, before performing the directional wavelet transform on the difference image f^{dir} . The balancing between curvelets and directional wavelets, L_{trans} is also flexible and can be tuned in our hybrid construction, depending on the importance of directional structure in the image or the computational time available.

This hybrid wavelet transform is implemented in the S2LET package.⁵ The current implementation is not optimised as it performs a full backward wavelet transform when only one scale needs to be transformed. This optimisation is left for future work.

²<http://www.s2let.org>

³<http://www.spinsht.org>

⁴<http://www.sothree.org>

⁵Support for hybrid wavelets in the S2LET package will be made public following the publication of this article.

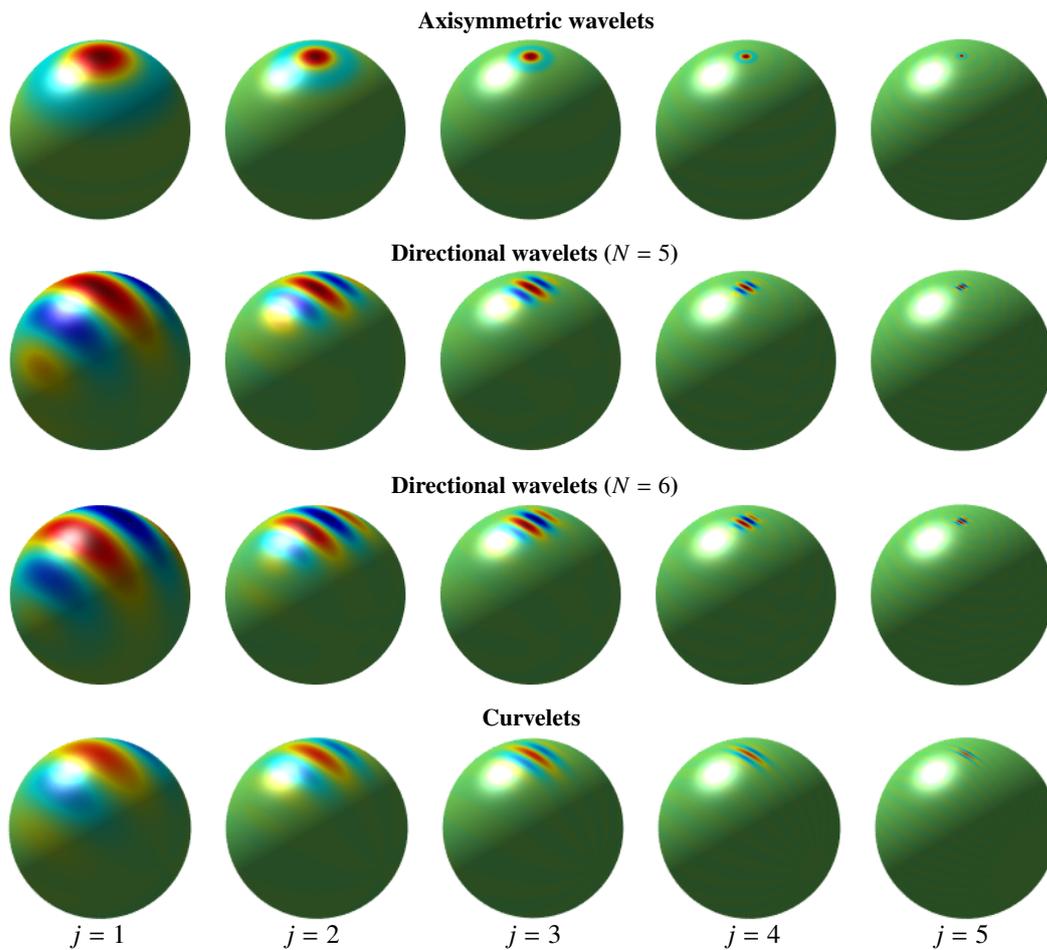


Figure 2: Scalar scale-discretised axisymmetric wavelets, directional wavelets ($N = 5$ and 6), and curvelets on the sphere for $L = 512$, from the first row to the fourth row.

2.3. Gradient operators on the sphere

The segmentation method developed herein requires the computation of gradients on the sphere $\nabla f = (\frac{\partial f}{\partial \theta}, \frac{\partial f}{\partial \phi})$, with the continuous magnitude of the gradient given by

$$\|\nabla f\| \equiv \sqrt{\left(\frac{\partial f}{\partial \theta}\right)^2 + \frac{1}{\sin^2 \theta} \left(\frac{\partial f}{\partial \phi}\right)^2}.$$

Discrete gradient operators for the equiangular sampling scheme adopted [38] are defined in [35]. The discrete magnitude of the gradient is simply defined by

$$\|\nabla f\| \equiv \sqrt{(\delta_\theta f)^2 + \frac{1}{\sin^2 \theta_t} (\delta_\phi f)^2}. \quad (5)$$

where δ_θ and δ_ϕ are finite difference operators. For more details of the discrete gradient operator please refer to [35].

2.4. Tight-frame based segmentation method

The following presents the generic tight-frame algorithm used in e.g. [7]:

$$f^{(i+\frac{1}{2})} = \mathcal{U}(f^{(i)}), \quad (6)$$

$$f^{(i+1)} = \mathcal{A}^\top \mathcal{T}_\lambda(\mathcal{A} f^{(i+\frac{1}{2})}), \quad i = 1, 2, \dots \quad (7)$$

Here \mathcal{A} and \mathcal{A}^\top are the tight-frame (wavelets in our case) forward and backward transforms respectively, $f^{(i)}$ is an approximate solution at the i -th iteration, \mathcal{U} is a problem-dependent operator (e.g. \mathcal{U} is the identity operator for a denoising problem), and $\mathcal{T}_\lambda(\cdot)$ is the soft-thresholding operator defined by

$$\mathcal{T}_\lambda(\vec{v}) \equiv [t_\lambda(v_1), \dots, t_\lambda(v_n)]^\top,$$

where $\vec{v} = [v_1, \dots, v_n]^\top \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}^+$ are a given vector and constant respectively, and

$$t_\lambda(v_k) \equiv \begin{cases} \text{sign}(v_k)(|v_k| - \lambda), & \text{if } |v_k| > \lambda, \\ 0, & \text{if } |v_k| \leq \lambda. \end{cases} \quad (8)$$

To obtain a binary result, where values 1 and 0 represent the object of interest and the background respectively, an iterative procedure was proposed in [9, 10] to gradually update an interval that contains pixel values of potential boundary pixels until the interval is empty. Note that the test image discussed in [9, 10] is assumed to have a low noise level and scaled to the range [0, 1].

The main segmentation procedures of [10] are as follows. Firstly, separate the given image to three parts by thresholding, i.e., area of background, area of object of interest, and the uncertainty-area which needs to be labelled as background or object in future steps. Secondly, denoise and smooth the uncertainty-area by the tight-frame algorithm [7] to get a new uncertainty-area which is smaller than the previous one. Thirdly, stop the algorithm when the uncertainty-area is empty (a binary result is then obtained), otherwise continue.

3. Spherical segmentation method

In [9, 10], the tight-frame based segmentation method is applied to Euclidean images but it is in principle extendable to a spherical domain and is compatible with different types of wavelets transforms. In this paper, based on the idea in the method [9, 10] on Euclidean images, we propose the wavelet-based segmentation framework on the sphere for spherical images.

The idea behind the method is to detect the candidates of possible pixels on (near) the boundary first, then gradually purify these boundary-like pixels via an iterative procedure until all pixels on the sphere are classified as inside or outside of a boundary. With the aid of the fact that possible pixels on the boundary have particular properties in terms

of pixel values and gradients, boundary-like pixels are detected and represented by a range $[a_0, b_0]$. Then, an iterative strategy shrinking this range is applied, so to keep removing pixels from it until the range itself is empty. All pixels are eventually classified either as in the foreground (the objects of interest) or in the background. Note that pixels in the foreground and in the background are represented respectively by values 1 and 0. When a binary result is obtained the algorithm stops.

The greater the anisotropic structure in the image, the more complicated the boundary-like pixels in $[a_0, b_0]$. Therefore, using an iterative procedure is particularly useful for images containing anisotropic structures. Otherwise, replacing the iterative procedure by thresholding is more economical (as demonstrated very effective in [11, 12, 13, 14]). In the following, we discuss each of the iterative steps of the proposed method in more detail.

Preprocessing. If f is contaminated with significant noise, a preprocessing step to suppress the noise is necessary. We use one iteration step of the tight-frame algorithm (7) to deal with the noise by soft thresholding, i.e.

$$\bar{f} = \mathcal{A}^\top \mathcal{T}_\lambda(\mathcal{A}f). \quad (9)$$

Note that \mathcal{A} here is a wavelet transform on sphere.

Initialisation. Let $\Lambda^{(0)}$ be the initial set of potential boundary pixels, which is identified by using the gradient of \bar{f} , i.e. pixels with gradient larger than a given threshold ϵ are in $\Lambda^{(0)}$, therefore

$$\Lambda^{(0)} \equiv \{k \in \mathbb{S}^2 \mid \|[\nabla \bar{f}]_k\|_1 > \epsilon\}. \quad (10)$$

Here $[\nabla \bar{f}]_k$ (cf. (5)) is the discrete gradient of \bar{f} at the k -th pixel on the sphere. Set $f^{(0)} = \bar{f}$, with $\Lambda^{(0)}$ defined in (10). We start the iterative process from $i = 0$. The i -th iteration is described in detail below.

Step 1: computing the range $[a_i, b_i]$. Given $\Lambda^{(i)}$, define $[a_i, b_i]$ by

$$a_i \equiv \max \left\{ \frac{\mu_-^{(i)} + \mu_-^{(i)}}{2}, 0 \right\}, \quad b_i \equiv \min \left\{ \frac{\mu_+^{(i)} + \mu_+^{(i)}}{2}, 1 \right\}, \quad (11)$$

where

$$\mu^{(i)} = \frac{1}{|\Lambda^{(i)}|} \sum_{k \in \Lambda^{(i)}} f_k^{(i)} \quad (12)$$

is the mean pixel value on $\Lambda^{(i)}$, $|\cdot|$ denotes the cardinality of the set, $f_k^{(i)}$ is the pixel value of pixel k in spherical image $f^{(i)}$, and $\mu_-^{(i)}$ and $\mu_+^{(i)}$ are defined by

$$\mu_-^{(i)} = \frac{\sum_{\{k \in \Lambda^{(i)} : f_k^{(i)} \leq \mu^{(i)}\}} f_k^{(i)}}{|\{k \in \Lambda^{(i)} : f_k^{(i)} \leq \mu^{(i)}\}|}, \quad \mu_+^{(i)} = \frac{\sum_{\{k \in \Lambda^{(i)} : f_k^{(i)} \geq \mu^{(i)}\}} f_k^{(i)}}{|\{k \in \Lambda^{(i)} : f_k^{(i)} \geq \mu^{(i)}\}|}. \quad (13)$$

Note that $\mu_-^{(i)}$ and $\mu_+^{(i)}$, the mean pixel values of the two sets separated by $\mu^{(i)}$, reflect the mean energies of the pixels on the boundary closer to the background and closer to the foreground respectively. The definition of $[a_i, b_i]$ in (11) is approximately half length of $[a_{i-1}, b_{i-1}]$, ensuring the shrinkage property of these ranges.

Step 2: thresholding the image into three parts. Using $[a_i, b_i] \subseteq [0, 1]$, we separate image $f^{(i)}$ into three parts — those below (set those pixel values that are smaller than a_i to 0), inside (stretch those pixel values between 0 and 1 using a simple linear contrast stretch), and above (set those pixel values that are larger than b_i to 1) the range, i.e.,

$$f_k^{(i+\frac{1}{2})} = \begin{cases} 0, & \text{if } f_k^{(i)} \leq a_i, \\ \frac{f_k^{(i)} - m_i}{M_i - m_i}, & \text{if } a_i \leq f_k^{(i)} \leq b_i, \\ 1, & \text{if } b_i \leq f_k^{(i)}, \end{cases} \quad \text{for all } k \in \mathbb{S}^2. \quad (14)$$

where

$$\begin{aligned} M_i &= \max\{f_k^{(i)} \mid a_i \leq f_k^{(i)} \leq b_i, k \in \Lambda^{(i)}\}, \\ m_i &= \min\{f_k^{(i)} \mid a_i \leq f_k^{(i)} \leq b_i, k \in \Lambda^{(i)}\}. \end{aligned} \quad (15)$$

The set of the remaining pixels that wait to be labelled is represented by

$$\Lambda^{(i+1)} = \{k \mid 0 < f_k^{(i+\frac{1}{2})} < 1, k \in \mathbb{S}^2\}. \quad (16)$$

Note that if $\Lambda^{(i+1)} = \emptyset$, the threshold image $f_k^{(i+\frac{1}{2})}$ is binary and the algorithm stops.

Remark 1. After obtaining $\Lambda^{(i+1)}$ from formula (16), the segmentation accuracy could be improved by making a correction of $\Lambda^{(i+1)}$, such as by adding the labelled but isolated (or wrongly labelled) pixels back to $\Lambda^{(i+1)}$ before moving to step (18) (or step (17)). We leave this to future work.

When the pixels in range $[a_i, b_i]$ can be classified easily (e.g. the number of pixels left to be classified is few therefore they are no longer critical to the final segmentation result) as the background or the objects of interest, the thresholding step for segmentation can be invoked:

$$f_k^{(i+\frac{3}{2})} = \begin{cases} 0, & \text{if } f_k^{(i+\frac{1}{2})} < \mu, \\ 1, & \text{if } f_k^{(i+\frac{1}{2})} \geq \mu, \end{cases} \quad \mu = \frac{1}{|\Lambda^{(i+1)}|} \sum_{k \in \Lambda^{(i+1)}} f_k^{(i)} \quad (17)$$

and the iteration terminates.

Step 3: spherical wavelets iteration. Let \mathcal{I} be the identity operator and $\mathcal{P}^{(i+1)}$ be the operator where the entry is 1 if the corresponding index is in $\Lambda^{(i+1)}$, and 0 otherwise. Then

$$f^{(i+1)} \equiv (\mathcal{I} - \mathcal{P}^{(i+1)})f^{(i+\frac{1}{2})} + \mathcal{P}^{(i+1)}\mathcal{A}^T\mathcal{T}_\lambda(\mathcal{A}f^{(i+\frac{1}{2})}). \quad (18)$$

Recall \mathcal{A} here represents a wavelet transform on sphere (e.g. axisymmetric wavelets, directional wavelets, curvelets, or hybrid wavelets). Note that the values of all pixels outside $\Lambda^{(i+1)}$ are either 0 or 1, hence the cost of (18) can be reduced significantly by applying the forward and backward wavelet transforms on pixels around $\Lambda^{(i+1)}$ only. This optimisation is left for future work.

Stopping criterion. As soon as all the pixels of $f^{(i+\frac{1}{2})}$ are either of value 0 or 1, or equivalently when $\Lambda^{(i)} = \emptyset$, the iteration is terminated, then all the pixels with value 1 constitute the objects of interest otherwise they are considered as background.

Algorithm 1 below summarises the steps required to segment a spherical image f by our segmentation method. Its convergence proof follows the proof given in [10]. In subsequent sections, algorithm 1 is referred to as WSSA for simplicity.

Algorithm 1: Wavelet-based Spherical Segmentation Algorithm (WSSA)

- 1 **Input:** given image $f \in L^2(\mathbb{S}^2)$
 - 2 Preprocessing by (9)
 - 3 Set $f^{(0)} = \bar{f}$ and $\Lambda^{(0)}$ by (10)
 - 4 **do**
 - 5 compute $[a_i, b_i]$ by (11)
 - 6 compute $f^{(i+\frac{1}{2})}$ by (14)
 - 7 stop if $f^{(i+\frac{1}{2})}$ is a binary image
 - 8 compute $\Lambda^{(i+1)}$ by (16)
 - 9 compute $f^{(i+1)}$ by (18) (or compute $f^{(i+\frac{3}{2})}$ by (17) then stop)
 - 10 $i = i + 1$
 - 11 **while** Stopping criterion is not reached;
-

A potential extension of our proposed spherical segmentation method is to interface with spherical neural networks, which have recently shown to be highly effective for analysing spherical images [20, 19, 22, 44, 26]. The first vital steps of it involve using a spherical neural network to effectively denoise and smooth the data on the sphere, i.e.

to replace the procedures shown in lines 2 and 9 of Algorithm 1. We leave the detailed implementation as our future work.

4. Experiments

We apply our method, namely WSSA, to various kinds of real-life images, including an Earth topographic map, light probe image, solar data-sets, and retina images projected on the sphere. Axisymmetric wavelets, directional wavelets, and hybrid wavelets constructed by combining the directional wavelets and curvelets are tested and their performances are compared. Algorithm 1 (WSSA) equipped with axisymmetric wavelets, directional wavelets and hybrid wavelets are referred to as WSSA-A, WSSA-D, and WSSA-H, respectively. The code to perform these spherical wavelets transforms used are available in the software package S2LET [28, 34]. The popular K-means method (e.g. [24, 27]) is implemented for comparison purposes. Here, the K-means method is applied to data on the sphere according to the pixels intensities, using the MATLAB built-in function `kmeans`. All the experiments are executed on a MacBook with 2.2 GHz Intel Core i7 processor and 16GB RAM.

Parameters. We set the spherical wavelet band-limit $L = 512$, the minimum angular scale $J_{\min} = 2$, and the number of directions probed by directional wavelets to be $N = 5$ and 6. We discretise the sphere \mathbb{S}^2 with size 512×1023 (refer to section 2.1), therefore $|\mathbb{S}^2| = 523776$. For the hybrid wavelets $L_{\text{trans}} \in \{32, 64\}$, which means that curvelets are used for bands up to $\ell \lesssim \{32, 64\}$ and directional wavelets for the remaining bands. Gaussian noise with standard deviation $\sigma = \|f\|_{\infty} 10^{-\text{SNR}/20}$ and 0 mean is added to the test data, where $\text{SNR} = 30$ dB and $\|\cdot\|_{\infty}$ is the infinity norm (referring to the maximum value). We fix the thresholding parameter $\bar{\lambda} = \sigma/4$ in (9) for denoising, and $\lambda = \sigma/100$ in (18) for segmentation during the spherical wavelet iterations.

4.1. Earth topographic map, light probe image and solar data-sets

Our first example is segmenting an Earth topographic map. The original Earth topography data are taken from the Earth Gravitational Model (EGM2008) publicly released by the U.S. National Geospatial-Intelligence Agency (NGA) EGM Development Team.⁶ The signal is band limited to $L = 512$ by performing a forward spherical harmonic transform, band limiting in harmonic space and transforming the signal back from its coefficients.

Table 1: Earth map in Fig. 3: Number of unclassified points at each iteration (i.e. $|\Lambda^{(i)}|$, where $\Lambda^{(i)}$ is the set of unclassified points and $|\cdot|$ denotes the cardinality of a set) and computation time in seconds. *The fourth and fifth columns represent the results of WSSA-D with $N = 5$ and 6, respectively.

	K-means	WSSA-A	WSSA-D*	WSSA-D*	WSSA-H
$ \mathbb{S}^2 $	523776	523776	523776	523776	523776
$ \Lambda^{(0)} $	-	111371	110373	111854	111184
$ \Lambda^{(1)} $	-	106977	104381	105222	106946
$ \Lambda^{(2)} $	-	25880	25938	26387	27681
$ \Lambda^{(3)} $	-	6352	6750	6645	6940
$ \Lambda^{(4)} $	-	1824	1995	1937	1972
$ \Lambda^{(5)} $	-	615	680	668	664
$ \Lambda^{(6)} $	-	229	247	269	254
$ \Lambda^{(7)} $	-	96	89	97	83
$ \Lambda^{(8)} $	-	28	26	38	27
$ \Lambda^{(9)} $	-	5	7	12	9
$ \Lambda^{(10)} $	-	0	2	2	0
$ \Lambda^{(11)} $	-	-	0	0	-
Time	< 1 s	51.9 s	200.5 s	217.2 s	883.5 s

⁶These data were downloaded and extracted using the tools available from Frederik Simons' webpage: <http://www.frederik.net>.

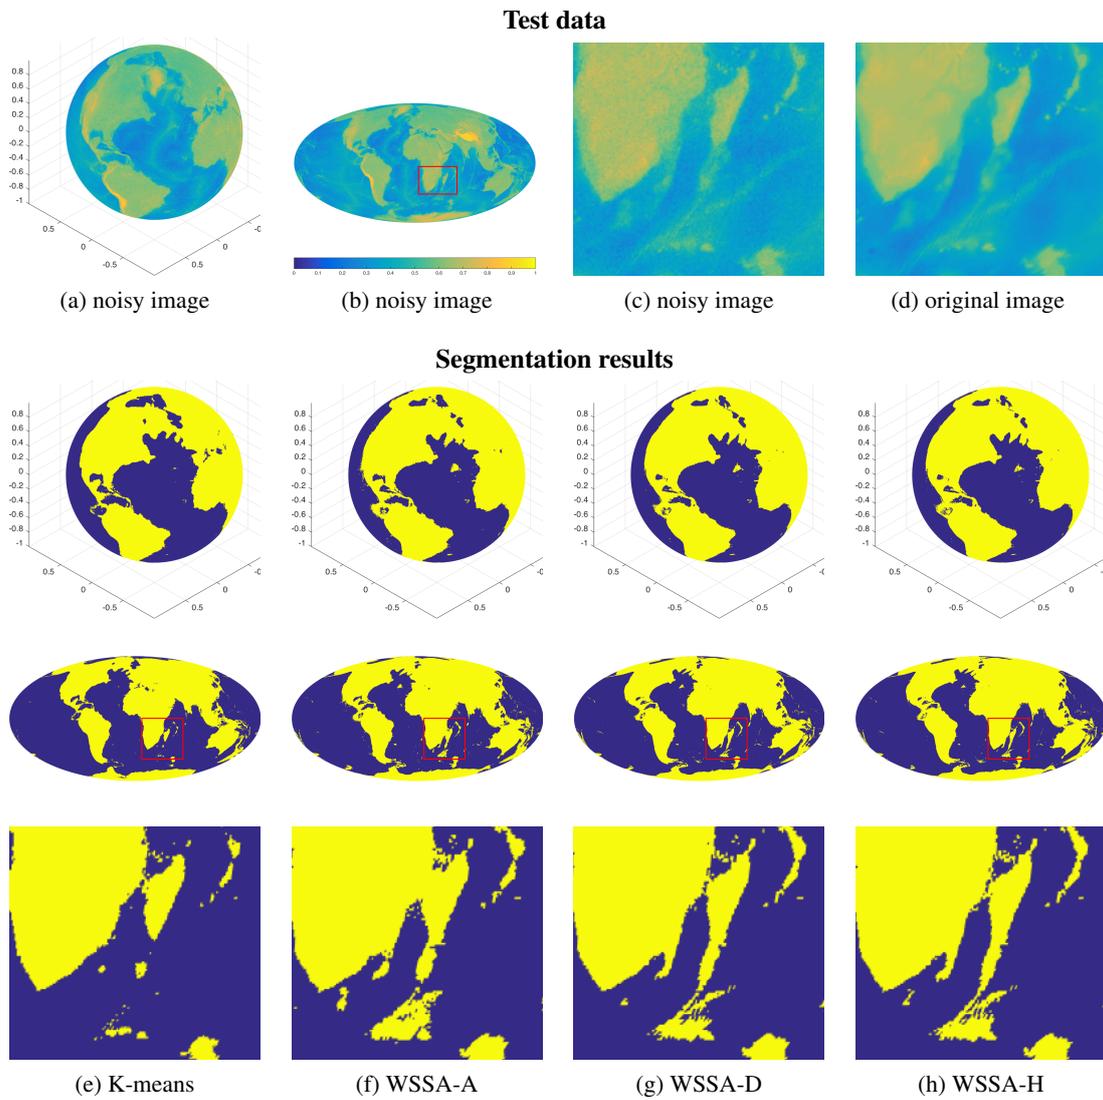


Figure 3: Results of the Earth topographic map. First row: noisy image shown on the sphere (a) and in 2D using a mollweide projection (b), and the zoomed-in red rectangle area of the noisy (c) and original images (d), respectively; Second to fourth rows from left to right: results of methods K-means (e), WSSA-A (f), WSSA-D (g) with $N = 5$ (odd N), and WSSA-H (h), respectively.

Fig. 3 shows the results of the K-means and our WSSA (-A, -D, and -H) method with $\epsilon = 0.02$ used in (10) to obtain the initial set $\Lambda^{(0)}$. Fig. 3 (a) and (b) are the test noisy image corrupted by Gaussian noise shown on the sphere and in 2D using a mollweide projection, respectively. Fig. 3 (c) shows the zoomed-in details of the red rectangle in Fig. 3 (b). For easy of comparison, Fig. 3 (d) shows the same zoomed-in area of Fig. 3 (c) with no noise added. The second to the fourth rows present the segmentation results shown on the sphere, in 2D using a mollweide projection, and with the zoomed-in details of the red rectangle area, respectively. From the results, we see that all the methods give reasonable segmentation results, i.e., the land and oceans are separated quite well (it should be noted that the separation is not necessarily strictly into land and sea as there is no requirement for the shore line to be the segmentation boundary). From the zoomed-in details, we see the WSSA method (Fig. 3 (f)–(h)) produces a better segmentation than the K-means method (Fig. 3 (e)). We also see that the WSSA method equipped with directional wavelets (Fig. 3 (g)) and hybrid wavelets (Fig. 3 (h)) are marginally better than using the axisymmetric wavelets (Fig. 3 (f)) in terms of preserving directional features in these data. WSSA-D and WSSA-H methods give very similar results (note that the hybrid method comprises the directional wavelets as a major component in this test).

Table 1 gives the number of unclassified points at each iteration ($|\Lambda^{(i)}|$) of the WSSA method and the computation time in seconds, from which we can see the WSSA method takes about 10 iterations to converge in general, where each iteration takes roughly the same amount of computation time since in each iteration the computation time is dominated by a complete round-trip of wavelet transforms (which has yet to be optimised utilising the data structure in the proposed segmentation method). Note that after the third iteration, the number of unclassified pixels is already very low compared with the whole number of pixels on the sphere $|\mathbb{S}^2|$. Moreover, from Table 1, we see the WSSA-H method needs the longest computation time, while the K-means needs the shortest; for the WSSA-D method, the greater N , the longer the computation time required.

Table 2: Light probe image - the Uffizi Gallery in Fig. 4: Number of unclassified points at each iteration and computation time in seconds. *The fourth and fifth columns represent the results of WSSA-D with $N = 5$ and 6, respectively.

	K-means	WSSA-A	WSSA-D*	WSSA-D*	WSSA-H
$ \mathbb{S}^2 $	523776	523776	523776	523776	523776
$ \Lambda^{(0)} $	-	35242	36372	36711	36534
$ \Lambda^{(1)} $	-	21246	21176	21430	21337
$ \Lambda^{(2)} $	-	5350	5371	5456	5437
$ \Lambda^{(3)} $	-	1453	1516	1491	1519
$ \Lambda^{(4)} $	-	402	434	404	431
$ \Lambda^{(5)} $	-	111	116	108	138
$ \Lambda^{(6)} $	-	30	29	29	30
$ \Lambda^{(7)} $	-	5	6	6	7
$ \Lambda^{(8)} $	-	0	0	0	0
Time	< 1 s	41.9 s	145.7 s	152.2 s	702.7 s

The second example is on segmenting a light probe image of a natural scene: the Uffizi Gallery in Florence⁷. A light probe image was created by taking two pictures of a mirrored ball ninety degrees apart and assembling the two radiance maps into a full sphere.

Fig. 4 shows the results of segmenting the light probe image of the Uffizi Gallery, with $\epsilon = 0.05$ used in the WSSA method. The same conclusions as those of the Earth map segmentation are obtained. Nonetheless, specific to the example here, in separating the sky and the bright parts of the windows within the test data the WSSA method is better at detecting detailed structures than the K-means method, for example, in the window frames shown in the zoomed-in figures in Fig. 4 (e)–(h). Again, the WSSA-D and WSSA-H methods are slightly better than the WSSA-A method which uses axisymmetric wavelets, but WSSA-A is faster (see Table 2).

The third example is the application of the segmentation algorithms to two different solar data-sets. Solar maps are

⁷ The data were downloaded from the webpage: <http://www.pauldebevec.com/Probes/>.

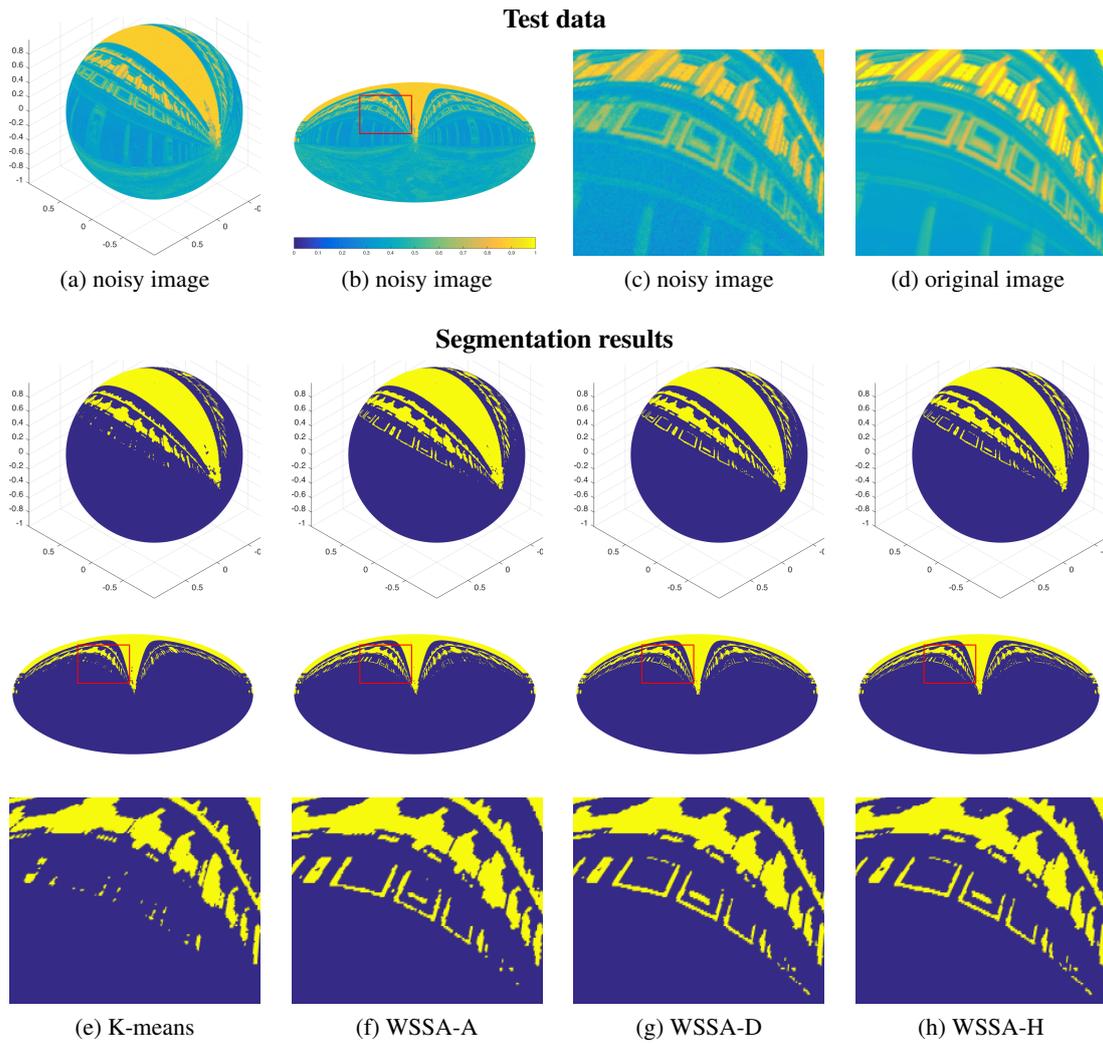


Figure 4: Results of light probe image - the Uffizi Gallery. First row: noisy image shown on the sphere (a) and in 2D using a mollweide projection (b), and the zoomed-in red rectangle area of the noisy (c) and original images (d), respectively; Second to fourth rows from left to right: results of methods K-means (e), WSSA-A (f), WSSA-D (g) with $N = 6$ (even N), and WSSA-H (h), respectively.

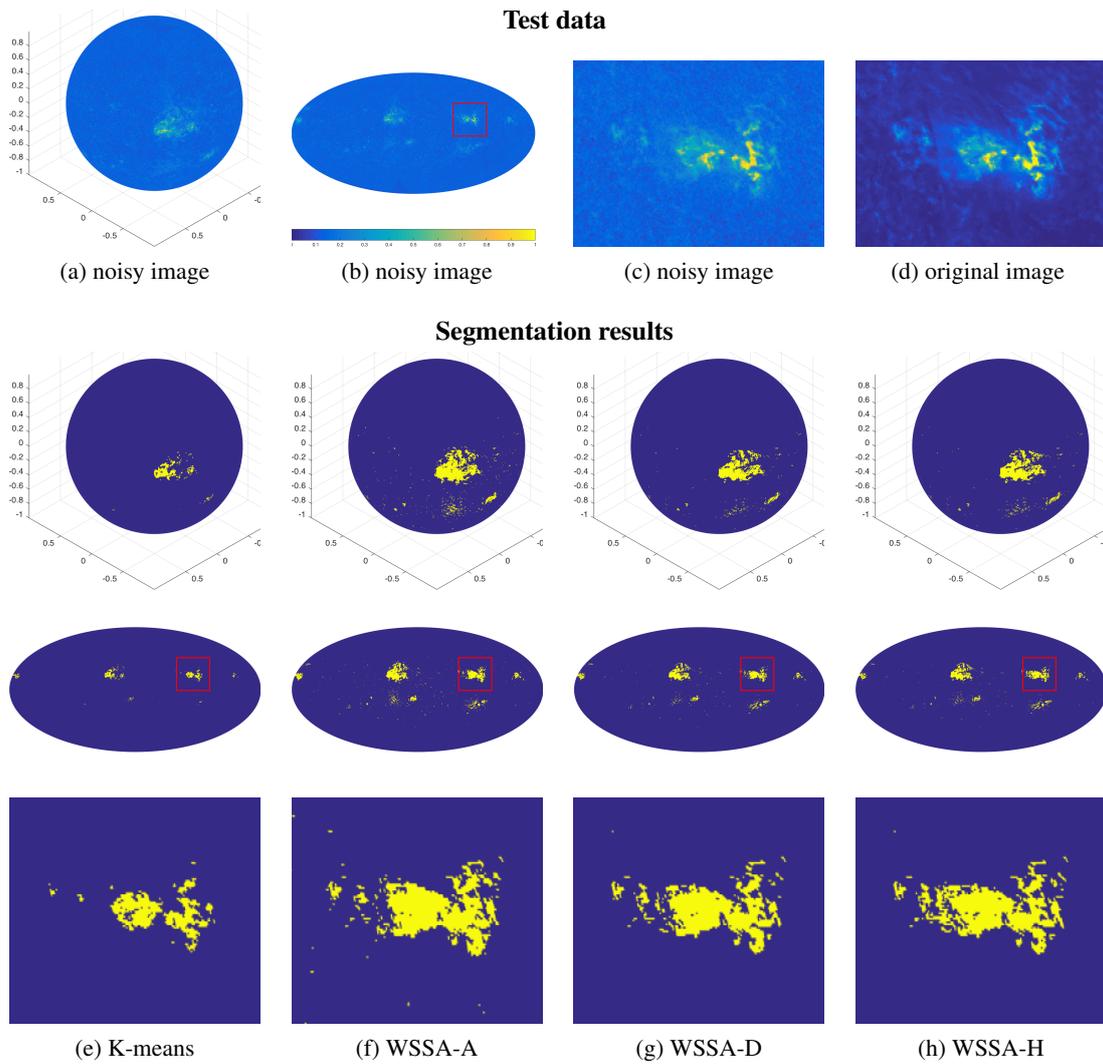


Figure 5: Results of solar map. First row: noisy image shown on the sphere (a) and in 2D using a mollweide projection (b), and the zoomed-in red rectangle area of the noisy (c) and original images (d), respectively; Second to fourth rows from left to right: results of methods K-means (e), WSSA-A (f), WSSA-D (g) with $N = 5$ (odd N), and WSSA-H (h), respectively.

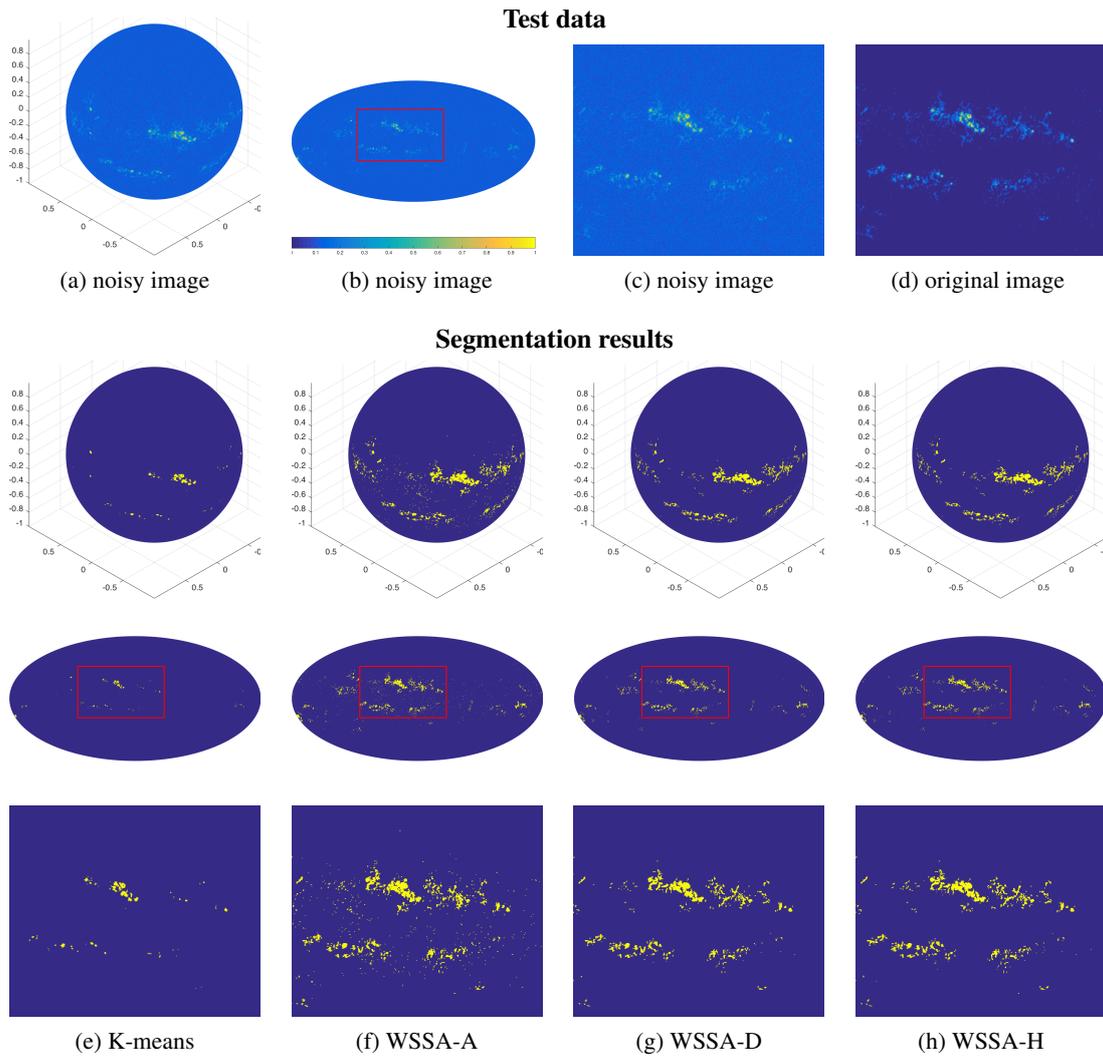


Figure 6: Results of solar map. First row: noisy image shown on the sphere (a) and in 2D using a mollweide projection (b), and the zoomed-in red rectangle area of the noisy (c) and original images (d), respectively; Second to fourth rows from left to right: results of methods K-means (e), WSSA-A (f), WSSA-D (g) with $N = 6$ (even N), and WSSA-H (h), respectively.

Table 3: Solar map in Fig. 5: Number of unclassified points at each iteration and computation time in seconds. *The fourth and fifth columns represent the results of WSSA-D with $N = 5$ and 6, respectively.

	K-means	WSSA-A	WSSA-D*	WSSA-D*	WSSA-H
$ \mathbb{S}^2 $	523776	523776	523776	523776	523776
$ \Lambda^{(0)} $	-	17324	14422	13960	13904
$ \Lambda^{(1)} $	-	14480	10827	10681	11326
$ \Lambda^{(2)} $	-	3468	2471	2484	2674
$ \Lambda^{(3)} $	-	884	644	649	714
$ \Lambda^{(4)} $	-	220	166	174	173
$ \Lambda^{(5)} $	-	60	45	43	36
$ \Lambda^{(6)} $	-	12	10	9	7
$ \Lambda^{(7)} $	-	0	0	1	1
$ \Lambda^{(8)} $	-	-	-	0	0
Time	< 1 s	34.1 s	124.3 s	151.8 s	682.2 s

very informative about solar activity, which have direct and indirect impacts on our activities on Earth. For illustrative purposes, we apply our method to segment two solar maps which show different solar features.

The first solar data, presented in Fig. 5, are obtained by synthetically sewing three spacecraft measurements taken on 8th July 2012 at wavelength 30.4 nm. The three instruments are SDO/AIA⁸, STEREO-A/SECCHI and STEREO-B/SECCHI⁹. The three spacecraft orbit around the Sun and together they give a full snapshot of the Sun in 360 degrees. The sewing procedures account for the angles each spacecraft instrument covered. Then by removing the overlapped observed regions between each pair of instruments and stitching the maps together, a complete snapshot of the Sun is obtained. Regions at solar latitudes affected by edge effects and the tilting of the observations have their intensity set to zero. A Gaussian filter was applied to smooth the resulting map. More sophisticated methods are required to properly combine and interpret the data from different solar instruments but this is beyond the scope of this article. Here we focus on segmenting the features seen in a snapshot of the Sun at the wavelength sensitive to solar flares. Parameter $\epsilon = 0.04$ is used in our WSSA method. From Fig. 5, one can see that WSSA-A, WSSA-D, and WSSA-H methods performed much better than the K-means method, preserving the directional features of the signal more completely. Among the three, as is seen from Fig. 5 (g) and (h), the quality of WSSA-H is slightly improved compared to the result from WSSA-D (in terms of the area of the sunspots), although the computation time is longer (see Table 3).

The second solar test data, presented in Fig. 6, is the radial-magnetic-field synoptic image of the Sun measured from 13th March 2001 to 9th April 2001 (Carrington Rotation 1974¹⁰), during which the solar activity peaked and a high number of sunspots were detected. The image shows the spatial variation of the strength of magnetic fields on the Sun (in our demonstration only absolute values are considered). We applied the K-means and the WSSA methods, with $\epsilon = 0.05$, to segment active magnetic regions on the Sun. As seen in Fig. 6 (e)-(h), the K-means method is able to pull out some of the sunspots (i.e. most magnetic-active regions) but it fails to capture the more diffusive and patchy features within the image. The WSSA-A, WSSA-D, and WSSA-H methods give quite similar results, which is mainly because the solar data themselves do not contain textures with strong directional information. At the same time, it can be seen that WSSA-D, and WSSA-H are more immune to noise.

4.2. Retina images on the sphere

In this example, we constructed test data containing very strong anisotropic structures from retina images. The retina images raised a very challenging segmentation problem because of the well known thin-vessel network (for example see [1, 25]). These kind of thin vessels are excellent for testing the ability of methods to tackle highly

⁸<http://sdo.gsfc.nasa.gov/>

⁹<http://www.stereo.rl.ac.uk/>

¹⁰<http://jsoc.stanford.edu/cgi-bin/hmisynop.pl?cr=1974&instrument=HMI&mag=Mag>

directional structures. The retina images tested here, (a) of Figs. 7 and 8, are from the DRIVE data-set¹¹ obtained from a diabetic retinopathy screening program in The Netherlands. They were acquired using a Canon CR5 non-mydratic 3CCD camera with a 45 degree field of view (FOV). Each image, captured using 8 bits per colour plane at 768×584 pixels, is scaled to $[0, 1]$ in our case. The FOV of each image is circular with a diameter of approximately 540 pixels. The images have been cropped around the FOV, and a mask image is provided that delineates the FOV.

To recover spherical retina images, for our tests, from the colour images (e.g. Fig. 7 (a1)), we transform the original planar images to the sphere in the following manner: (1) extract the green channel from the colour image, i.e. extract Fig. 7 (b1) from (a1); (2) obtain the background of the green channel by implementing the MATLAB built-in function `medfilt2` then remove it from Fig. 7 (b1) to get figure (c1); (3) add Gaussian noise to (c1) to generate the noisy image (d1) in Fig. 7; (4) project Fig. 7 (d1) to the spherical coordinate system to form the spherical retina image as our test data, shown in Fig. 7 (a) and (b). Fig. 7 (c) and (d) are the zoomed-in details of the rectangle in figure (b1) and in the projected image of figure (c1), respectively. Fig. 8 is generated and arranged using the same way as that in Fig. 7.

For the generated spherical retina images, (b) of Figs. 7 and 8, we use $\epsilon = 0.04$ in our WSSA method. The third to the fifth rows of the figures show the segmentation results. Clearly, the K-means method failed to identify most of the vessels (see the first column), while our WSSA method detected most of the vessels on the sphere (see the second until the fourth columns). After comparing the zoomed-in results, (f) – (h) of Figs. 7 and 8, we conclude that the WSSA-D and WSSA-H methods give better results than the WSSA-A method, the results of which contain more non-vessel artefacts. From Fig. 7, we see the improvement in the result of WSSA-H compared with the result of WSSA-D in terms of suppressing those non-vessel components about the north pole. Table 4 presents the time performance of each method, which is consistent with the conclusions obtained in the previous examples, i.e., the more directional the wavelet transform, the longer the computation time required. In particular, to test the property of the WSSA-H method, we consider the hybrid wavelets with $L_{\text{trans}} = 32$ and $L_{\text{trans}} = 64$ in Fig. 8 and Fig. 7, respectively, meaning that bands $\ell \lesssim 32$ and $\ell \lesssim 64$ are probed by curvelets, respectively, and the other bands are probed by directional wavelets. Their performance are given in Table 4. We see that the higher the L_{trans} , the longer computation time needed in the WSSA-H method, while the improvement in segmentation is small for these cases.

Table 4: Retina image in Fig. 8: Number of unclassified points at each iteration and computation time in seconds. *The fourth and fifth columns represent the results of WSSA-D with $N = 5$ and 6, respectively. †The last two columns represent the results of WSSA-H with $L_{\text{trans}} = 32$ and 64 for curvelets, respectively.

	K-means	WSSA-A	WSSA-D*	WSSA-D*	WSSA-H†	WSSA-H†
$ \bar{S}^2 $	523776	523776	523776	523776	523776	523776
$ \Lambda^{(0)} $	-	1640	1914	1934	1934	1928
$ \Lambda^{(1)} $	-	46611	28376	49958	12237	12088
$ \Lambda^{(2)} $	-	11366	6457	16346	3040	2945
$ \Lambda^{(3)} $	-	3095	1783	5001	852	832
$ \Lambda^{(4)} $	-	949	512	1603	257	258
$ \Lambda^{(5)} $	-	301	175	538	79	76
$ \Lambda^{(6)} $	-	95	63	183	19	25
$ \Lambda^{(7)} $	-	31	20	49	1	5
$ \Lambda^{(8)} $	-	10	6	18	0	0
$ \Lambda^{(9)} $	-	4	0	5	-	-
$ \Lambda^{(10)} $	-	0	-	0	-	-
Time	< 1 s	50.66 s	160.54 s	197.0 s	789.6 s	4538.9 s

¹¹<http://www.isi.uu.nl/Research/Databases/DRIVE/>

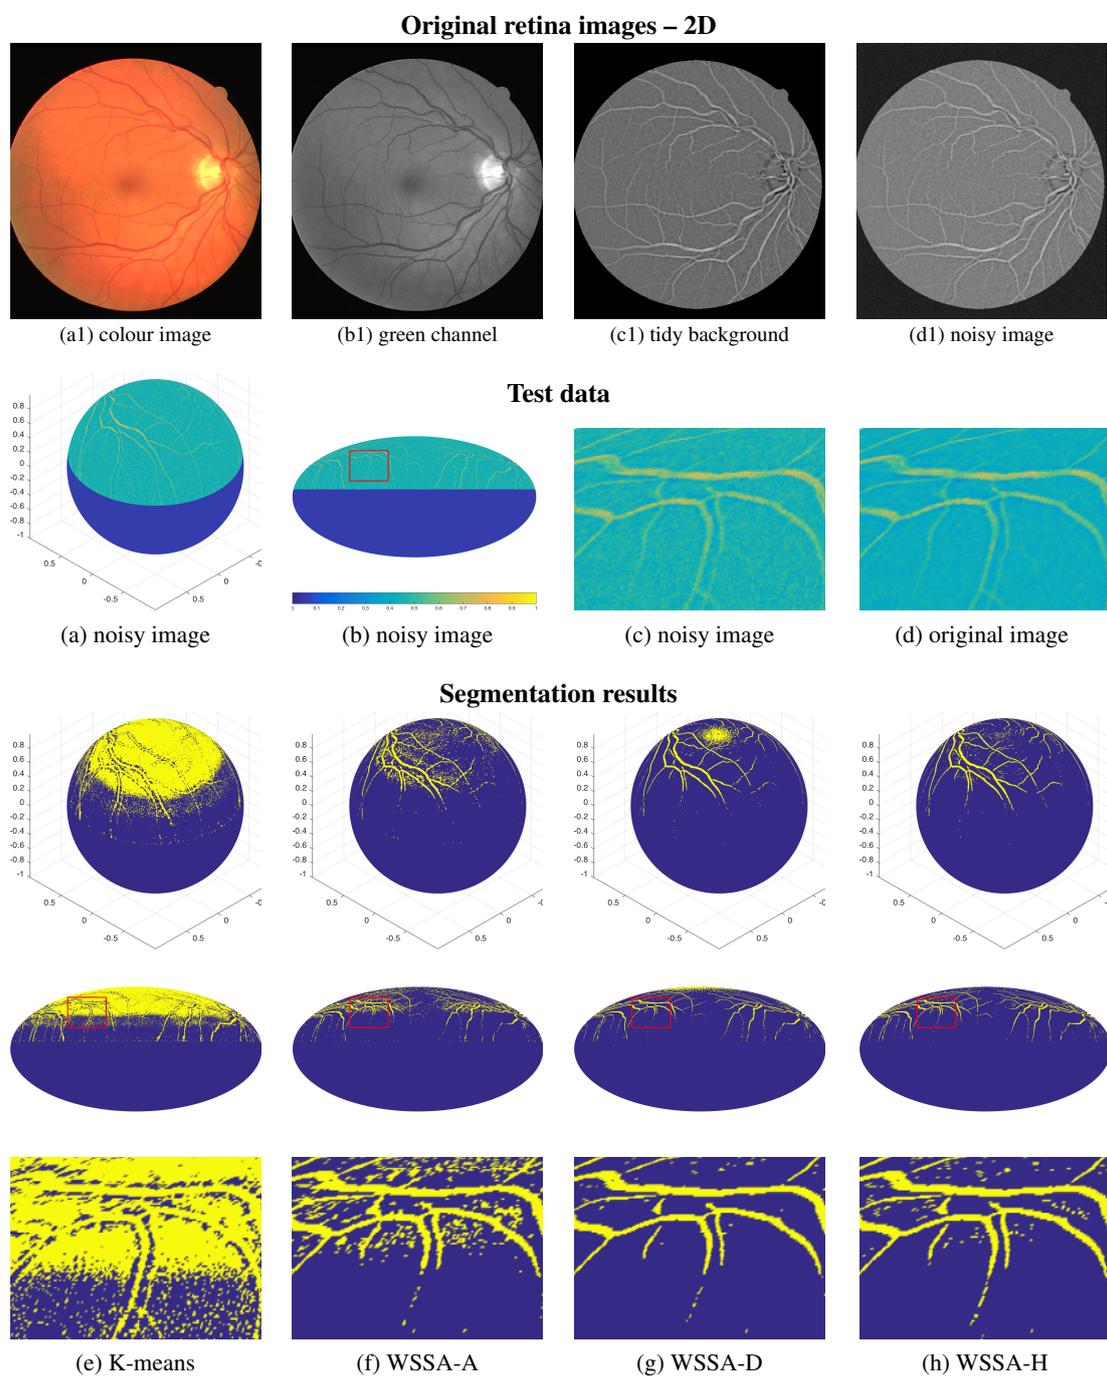


Figure 7: Results of retina image. First row: original retina image (a1), its green channel (b1), the image after removing its background (c1) and the noisy image (d1) respectively; Second row: projected retina image shown on the sphere (a) and in 2D using a mollweide projection (b), and the zoomed-in red rectangle area of the noisy (c) and original images (d), respectively; Third row to fifth row: results of methods K-means (e), WSSA-A (f), WSSA-D (g) with $N = 5$ (odd N), WSSA-H (h) with $L_{\text{trans}} = 64$ for curvelets, respectively.

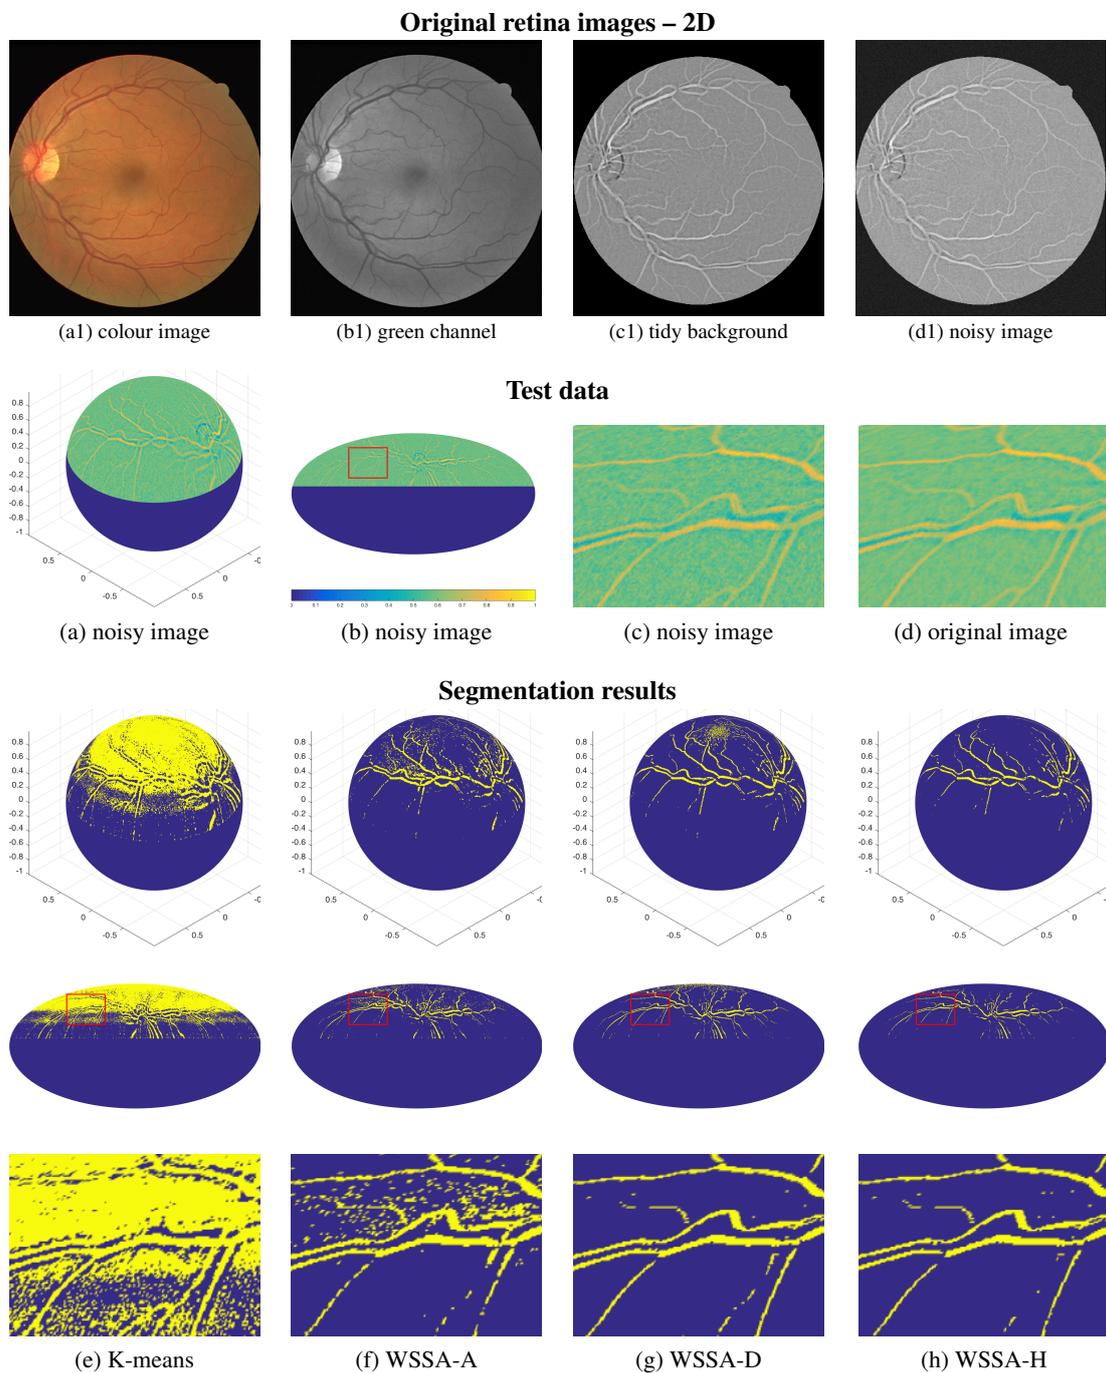


Figure 8: Results of retina image. First row: original retina image (a1), its green channel (b1), the image after removing its background (c1) and the noisy image (d1) respectively; Second row: projected retina image shown on the sphere (a) and in 2D using a mollweide projection (b), and the zoomed-in red rectangle area of the noisy (c) and original images (d), respectively; Third row to fifth row: results of methods K-means (e), WSSA-A (f), WSSA-D (g) with $N = 6$ (even N), WSSA-H (h) with $L_{\text{trans}} = 32$ for curvelets, respectively.

5. Conclusions

In this paper we proposed a wavelet-based segmentation method (WSSA) for spherical images, which is, to the best of our knowledge, the first method performing segmentation directly on the sphere. The method is compatible with any invertible wavelet transform constructed on the sphere (e.g. axisymmetric wavelets, directional wavelets, curvelets, or hybrid wavelets). Consequently, WSSA is very flexible and can be equipped with spherical wavelets appropriate for the texture property of the given spherical data of interest. WSSA needs just a few iterations to converge, and the main computation within each iteration is the pair of forward and backward wavelet transforms. We applied our WSSA method to several real-world problems, i.e., the Earth topographic map, a light probe image, two Solar maps, and two projected spherical retina images. The comparisons with the K-means method and different types of wavelets demonstrate that the WSSA method is an efficient and effective spherical segmentation method and is superior to K-means. One important future work will be focusing on purifying the uncertainty area $\Lambda^{(i)}$ at each step in each iteration of WSSA to improve segmentation quality according to specific applications. Another important future work will be extending the proposed spherical segmentation method by exploiting the recent developments of spherical neural networks.

Acknowledgements

This work is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) by grant EP/M011852/1 and EP/M011089/1. We thank Professor Raymond H. F. Chan in City University, Hong Kong, for the very helpful discussion. We also thank Dr David Peres-Suarez for his help in providing the first set of solar data. We would also like to thank the editors and anonymous reviewers for their valuable comments and suggestions to improve the quality of the paper.

References

- [1] D. Adjeroh, U. Kandaswamy, and O. Vernon. Texton-based segmentation of retinal vessels. *JOSA A*, 24(5):1384–1393, 2007.
- [2] J. Antoine, L. Demanet, L. Jacques, and P. Vandergheynst. Wavelets on the sphere: implementation and approximations. *Appl. Comput. Harmon. Anal.*, 13:177–200, 2002.
- [3] P. Audet. Toward mapping the effective elastic thickness of planetary lithospheres from a spherical wavelet analysis of gravity and topography. *Phys. Earth Planet In.*, 226(0):48–82, 2014.
- [4] S. Arivazhagan and L. Ganesan. Texture segmentation using wavelet transform. *Pattern Recognition Letters*, 24, 3197–3203, 2003.
- [5] P. Baldi, G. Kerkycharian, D. Marinucci, and D. Picard. Asymptotics for spherical needlets. *Ann. Stat.*, 37(3):1150–1171, 2009.
- [6] L. Bar, T. F. Chan, G. Chung, M. Jung, N. Kiryati, N. Sochen and L. A. Vese. Mumford and Shah Model and Its Applications to Image Segmentation and Image Restoration. *Handbook of Mathematical Methods in Imaging*, 2nd edition, Editor: O. Scherzer, Springer 2015 (online).
- [7] J. Cai, R. Chan, and Z. Shen. A framelet-based image inpainting algorithm. *Appl. Comput. Harmon. Anal.*, 24, 131–149, 2008.
- [8] X. Cai. Variational image segmentation model coupled with image restoration achievements. *Pattern Recognition*, 48:2029–2042, 2015.
- [9] X. Cai, R. Chan, S. Morigi, and F. Sgallari. Framelet-based algorithm for segmentation of tubular structures. *SSVM 2011, LNCS6667. Springer*, 2011.
- [10] X. Cai, R. Chan, S. Morigi, and F. Sgallari. Vessel segmentation in medical imaging using a tight-frame based algorithm. *SIAM J Imaging Sci.*, 6(1):464–486, 2013.
- [11] X. Cai, R. Chan, M. Nikolova, and T. Zeng. A three-stage approach for segmenting degraded color images: smoothing, lifting and thresholding (SLaT). *J Sci Comput.*, 72:1313–1332, 2017.
- [12] X. Cai, R. Chan, C.-B. Schönlieb, G. Steidl, and T. Zeng. Linkage between Piecewise Constant Mumford-Shah model and ROF model and Its Virtue in Image Segmentation. *arXiv:1807.10194*, 2018.
- [13] X. Cai, R. Chan, and T. Zeng. A two-stage image segmentation method using a convex variant of the Mumford-Shah model and thresholding. *SIAM J Imaging Sci.*, 6(1):368–390, 2013.
- [14] X. Cai and G. Steidl. Multiclass segmentation by iterated rof thresholding. In *EMMCVPR, LNCS, Springer*, 237–250, 2013.
- [15] E. Candés and D. Donoho. Continuous curvelet transform: I. resolution of the wavefront set. *Appl. Comput. Harmon. Anal.*, 19(2):162–197, 2005.
- [16] E. Candés and D. Donoho. Continuous curvelet transform: II. discretization and frames. *Appl. Comput. Harmon. Anal.*, 19(2):198–222, 2005.
- [17] T. Chan and L. Vese. Active contours without edges. *IEEE Trans. Image Process.*, 10:266–277, 2001.
- [18] J. Chan, B. Leistedt, T. D. Kitching, and J. D. McEwen. Second-generation curvelets on the sphere. *IEEE Trans. Sig. Proc.*, 65(1): 5–14, 2017.
- [19] T. Cohen, M. Geiger, J. Koehler, and M. Welling. Spherical CNNs. *Proc. of ICLR*, 2018.
- [20] B. Coors, A. Condurache, and A. Geiger. SphereNet: learning spherical representations for detection and classification in omnidirectional images. *ECCV*, 2018.

- [21] M. Do and M. Vetterli. The contourlet transform: an efficient directional multiresolution image representation. *IEEE Trans. Image Process.*, 14(12):2091–2106, 2004.
- [22] C. Esteves, C. Allen-Blanchette, A. Makadia, and K. Daniilidis. Learning SO(3) equivariant representations with spherical CNNs. *ECCV*, 2018.
- [23] E. Franchini, S. Morigi, and F. Sgallari. Segmentation of 3D tubular structures by a PDE-based anisotropic diffusion model. M. Dæhlen et al. (eds.): MMCS 2008, LNCS5862, pp. 224–241, 2010, *Springer-Verlag Berlin Heidelberg*, 2010.
- [24] J. HARTIGAN AND M. WANG, *A K-means clustering algorithm*, Applied Statistics, 28(1979), pp. 100–108.
- [25] H. Jelinek, M. Cree, J. Leandro, J. Soares, R. Cesar, and A. Luckie. Automated segmentation of retinal blood vessels and identification of proliferative diabetic retinopathy. *JOSA A*, 24(5):1448–1456, 2007.
- [26] C. Jiang, J. Huang, K. Kashinath, Prabhat, P. Marcus, and M. Niessner. Spherical CNNs on Unstructured Grids. *ICLR*, 2019.
- [27] T. KANUNGO, D. MOUNT, N. NETANYAHU, C. PIATKO, R. SILVERMAN, AND A. WU. An efficient k-means clustering algorithm: Analysis and implementation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24:881–892, 2002.
- [28] B. Leistedt, J. D. McEwen, P. Vandergheynst, and Y. Wiaux. S2LET: A code to perform fast wavelet analysis on the sphere. *Astronomy & Astrophysics*, 558 (A128):1–9 2013.
- [29] S. Li and Y. Hai. A full-view spherical image format. *ICPR*, pp. 2337–2340, 2010.
- [30] L. Lorigo, O. Faugeras, E. Grimson, et al. Curves: curve evolution for vessel segmentation. *Medical Image Analysis*, 5, 195–206, 2001.
- [31] J. D. McEwen, M. Büttner, B. Leistedt, H. V. Peiris, and Y. Wiaux. A novel sampling theorem on the rotation group. *IEEE Trans. Signal Process. Lett.*, 22(12):2425–2429, 2015.
- [32] J. D. McEwen, C. Durastanti, and Y. Wiaux. Localisation of directional scale-discretised wavelets on the sphere. *Applied Comput. Harm. Anal.*, 44(1):59–88, 2018.
- [33] J. D. McEwen, M. P. Hobson, D. J. Mortlock and A. N. Lasenby. Fast directional continuous spherical wavelet transform algorithms. *IEEE Trans. Signal Process.*, 55(2):520–529, 2007.
- [34] J. D. McEwen, B. Leistedt, M. Büttner, H. V. Peiris, and Y. Wiaux. Directional spin wavelets on the sphere. *IEEE Trans. Signal Process.*, arXiv:1509.06749, 2015.
- [35] J. D. McEwen, G. Puy, J. Thiran, P. Vandergheynst, D. Ville, and Y. Wiaux. Sparse image reconstruction on the sphere: implications of a new sampling theorem. *IEEE Trans. Image Process.*, 22(6):2275–2285, 2013.
- [36] J. D. McEwen, D. Vandergheynst, Y. Wiaux. On the computation of directional scale-discretized wavelet transforms on the sphere. *SPIE Wavelets and Sparsity XV*, Vol. 8858, 2013.
- [37] J. D. McEwen, P. Vielva, Y. Wiaux., R. B. Barreiro, L. Cayón, M. P. Hobson, A. N. Lasenby, E. Martínez-González, and J. L. Sanz. Cosmological applications of a wavelet analysis on the sphere. *J. Fourier Anal. and Appl.*, 13(4):495–510, 2007.
- [38] J. D. McEwen and Y. Wiaux. A novel sampling theorem on the sphere. *IEEE Trans. Sig. Proc.*, 59(12):5876–5887, 2011.
- [39] O. Michailovich and Y. Rathi. On approximation of orientation distributions by means of spherical ridgelets. *IEEE Trans. Sig. Proc.*, 19(2):461–477, 2010.
- [40] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Commun. Pur. Appl. Math.*, 42:577–685, 1989.
- [41] D. Nain, S. Haker, A. Bobick, and A. Tannenbaum. Multiscale 3-D shape representation and segmentation using spherical wavelets. *IEEE Trans. Medical Imaging*, 26(4):598–618, 2007.
- [42] F. Narcowich, P. Petrushev and D. Ward. Localized tight frames on spheres. *SIAM J. Math. Anal.*, 38(2):574–594, 2006.
- [43] M. Nikolova. A variational approach to remove outliers and impulse noise. *J. Math. Imag. Vis.*, 20(1-2): 99–120, 2004.
- [44] N. Perraudin, M. Defferrard, T. Kacprzak, and R. Sgier. DeepSphere: efficient spherical convolutional neural network with HEALPix sampling for cosmological applications. *arXiv:1810.12186*, 2019.
- [45] Y. Rathi, O. Michailovich, K. Setsompop, S. Bouix, M. E. Shenton, and C.-F. Westin. Sparse multi-shell diffusion imaging. *Medical image computing and computer-assisted intervention: MICCAI, International Conference on Medical Image Computing and Computer-Assisted Intervention*, 14(2):58–65, 2011.
- [46] T. Riklin-Raviv, N. Sochen, and N. Kiryati. On symmetry, perspectivity and level-set based segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 31:1458–1471, 2009.
- [47] A. Ron and Z. Shen. Affine systems in $L_2(\mathbb{R}^d)$: the analysis of the analysis operator. *J. Funct. Anal.*, 148, 408–447, 1997.
- [48] O. Ronneberger, Q. Wang, and H. Burkhardt. Fast and robust segmentation of spherical particles in volumetric data sets from bright field microscopy. *In Proc. of the ISBI*, pp. 372–375, 2008.
- [49] L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D.*, 60(1-4): 259–268, 1992.
- [50] J. Schmitt, J. L. Starck, J. M. Casandjian, J. Fadili, and I. Grenier. Multichannel Poisson denoising and deconvolution on the sphere: application to the Fermi Gamma-ray Space Telescope. *Astronomy & Astrophysics*, 546:A114, 2012.
- [51] J. Shi and J. Malik. Normalized Cuts and Image Segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(8):888–905, 2000.
- [52] F. Simons, I. Loris, G. Nolet, I. C. Daubechies, S. Voronin, J. S. Judd, P. A. Vetter, J. Charléty, and C. Vonesch. Solving or resolving global tomographic models with spherical wavelets, and the scale and sparsity of seismic heterogeneity. *Geophys. J. Int.*, 187:969–988, 2011.
- [53] M. Storath and A. Weinmann. Fast partitioning of vector-valued images. *SIAM J Imaging Sci.*, 7(3):1826–1852, 2014.
- [54] J. Starck, Y. Moudden, P. Abrial and M. Nguyen. Wavelets, ridgelets and curvelets on the sphere. *Astronomy & Astrophysics*, 446(3):1191–1204, 2006.
- [55] C. Wallis, Y. Wiaux, J. D. McEwen. Sparse image reconstruction on the sphere: analysis and synthesis. *IEEE Trans. Image Process.*, 26(11):5176–5187, 2017.
- [56] Y. Wiaux, J. D. McEwen, P. Vandergheynst, and O. Blanc. Exact reconstruction with directional wavelets on the sphere. *Mon. Not. Roy. Astron. Soc.*, 388(2):770–788, 2008.