Rotation Invariance and Directional Sensitivity: Spherical Harmonics versus Radiomics Features

Adrien Depeursinge[1,2(✉)], Julien Fageot[1], Vincent Andrearczyk[2], John Paul Ward[3], and Michael Unser[1]

[1] Biomedical Imaging Group, Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

[2] Institute of Information Systems, University of Applied Sciences Western Switzerland (HES-SO), Sierre, Switzerland
adrien.depeursinge@hevs.ch

[3] Department of Mathematics, North Carolina A&T State University, Greensboro, NC, USA

Abstract. We define and investigate the Local Rotation Invariance (LRI) and Directional Sensitivity (DS) of radiomics features. Most of the classical features cannot combine the two properties, which are antagonist in simple designs. We propose texture operators based on spherical harmonic wavelets (SHW) invariants and show that they are both LRI and DS. An experimental comparison of SHW and popular radiomics operators for classifying 3D textures reveals the importance of combining the two properties for optimal pattern characterization.

Keywords: Radiomics · 3D texture · Spherical harmonics · Wavelets

1 Introduction

Radiomics aims at establishing potential links between radiological image content and disease phenotypes [1]. Its potential for personalized medicine is enormous, thanks to the ability to accurately identify disease subtypes and responders to therapy in a noninvasive fashion. Its success fully relies on the relevance of extracted quantitative image measures for characterizing the manifestation of disease under the form of local tissue alterations (i.e. 3D texture) observed in Computed Tomography (CT), Magnetic Resonance Imaging (MRI) and Positron Emission Tomography (PET). The tissue alterations can be diverse, including necrosis, angiogenesis, fibrosis, cell proliferation (e.g. densification and higher metabolism) [2]. The latter induce corresponding imaging signatures in terms of 3D low-level patterns (e.g. blobs, intersecting surfaces and curves, Fig. 1). Such patterns are characterized by discriminative directional properties and have arbitrary 3D orientations. As a consequence, optimal radiomics image operators must combine Local Rotation Invariance (LRI) with Directional Sensitivity (DS) [3].

© Springer Nature Switzerland AG 2018
Y. Shi et al. (Eds.): MLMI 2018, LNCS 11046, pp. 107–115, 2018.
https://doi.org/10.1007/978-3-030-00919-9_13

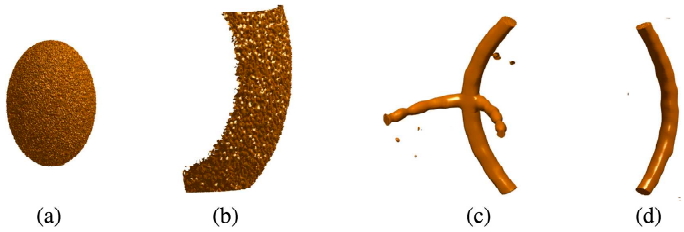

Fig. 1. 3D low-level patterns associated with biomedical tissue architectures. The latter include ellipsoid blobs (a), tumor or organ walls (b), and diverse vascular configurations (c), (d). These patterns can have arbitrary orientations in 3D medical images.

In this paper, we first formally define the notion of LRI and discuss how it relates to the DS of operators. We then propose a novel texture operator based on Spherical Harmonic Wavelet (SHW) invariants that can combine the two properties. Second, we qualitatively compare SHW invariants with classic 3D radiomics features in terms of LRI and DS. Finally, we evaluate experimentally the importance of the latter properties for 3D texture classification.

2 Materials and Methods

2.1 Notations

Vectors are written with bold symbols. The Fourier transform of an integrable function $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ is $\widehat{f}(\boldsymbol{\omega}) = \int_{\mathbb{R}^3} f(\boldsymbol{x}) \mathrm{e}^{\mathrm{j}\boldsymbol{\omega}^T \boldsymbol{x}} \mathrm{d}\boldsymbol{x}$. It is naturally extended to any square-integrable function in the usual manner. We use spherical coordinates (r, θ, ϕ) in Fourier domain, with radial distance $r \geq 0$, polar angle $\phi \in [0, 2\pi)$, and azimuthal angle $\theta \in [0, \pi)$.

2.2 Local Rotation Invariance and Directional Sensitivity

We follow the framework introduced in [3], where it is proposed that any texture analysis approach can be decomposed into (i) the application of a *texture operator* $\mathcal{G}$ associating to an image $f(\boldsymbol{x})$ a *response map* $h(\boldsymbol{x}) = \mathcal{G}\{f\}(\boldsymbol{x})$ and (ii) the aggregation of the response map information, transforming $h(\boldsymbol{x})$ into a scalar measurement through the use of an *aggregation function*[1] on a Volume Of Interest (VOI). This approach allows focusing on the properties of the texture operator in terms of LRI and DS, which we define as follows.

Definition 1. *A texture operator $\mathcal{G}$ is* Locally Rotation Invariant (LRI) *if the response map at $\boldsymbol{x}_0$ is not affected by the rotation of the input image around $\boldsymbol{x}_0$. In other terms, if, for any rotation operator $\mathcal{R}$, any $\boldsymbol{x}_0 \in \mathbb{R}^3$, and any input image $f \in L_2(\mathbb{R}^3)$, we have*

$$\mathcal{G}\{f\}(\boldsymbol{x}_0) = \mathcal{G}\{\mathcal{R}_{\boldsymbol{x}_0} f\}(\boldsymbol{x}_0), \tag{1}$$

[1] It is worth noting that steps (i) and (ii) are repeated multiple times in Convolutional Neural Networks (CNN).

where $\mathcal{R}_{\boldsymbol{x}_0}$ is a translated version of the rotation operator $\mathcal{R}$ that is centered around $\boldsymbol{x}_0$ instead of the origin $\mathbf{0}$ as $\mathcal{R}_{\boldsymbol{x}_0} = \mathcal{T}_{\boldsymbol{x}_0}\mathcal{R}\mathcal{T}_{-\boldsymbol{x}_0}$, with $\mathcal{T}_{\boldsymbol{x}_0}$ is a translation operator by $\boldsymbol{x}_0$.

The LRI is highly desirable for texture analysis, but is antagonist with the will of being sensitive to directional features to avoid mixing blobs, edges and ridges. For instance, it can be shown that a convolutional texture operator of the form $\mathcal{G}\{f\} = g * f$ is LRI if and only if the filter g is isotropic, therefore insensitive to the directional features of the input signal f. It follows that operators combining LRI and DS require using more complex designs such as SHW invariants presented in the next section.

2.3 Spherical Harmonic Wavelets

We design operators combining LRI and DS based on SHW. We use wavelets to characterize the local image information at different scales. The wavelet decomposition is computed efficiently due to the filterbank structure [4]. In what follows, we denote by $\widehat{g}(r)$ the 3-dimensional isotropic wavelet of interest, that only depends on the radius r. In practice, we use the Meyer wavelet [4] known for its good localization in space and frequency. The scale index is $i \in \mathbb{Z}$ and we set $\widehat{g}_i(r) = 2^{2i/3}\widehat{g}(2^i r)$ the wavelet at different scale (the normalization is chosen so that each g_i has norm 1). In what follows, $\mathcal{I} = \{i_{\min}, \dots, i_{\max}\}$ always refer to a finite subset of $\mathbb{Z}$.

The directional information is taken into account using spherical harmonics. The family of spherical harmonics is denoted by $(Y_n^m)_{n\geq 0, m\in\{-n,\dots,n\}}$, where n is called the degree and m the order of Y_n^m. Spherical harmonics form an orthonormal basis for square-integrable functions in the $2D$-sphere $\mathbb{S}^2$. They are defined as [5]

$$Y_n^m(\theta,\phi) = A_n^m P_n^{|m|}(\cos(\theta))\mathrm{e}^{\mathrm{j}m\phi}, \tag{2}$$

with $A_n^m = (-1)^{(m+|m|)/2}\left(\frac{2n+1}{4\pi}\frac{(n-|m|)!}{(n+|m|)!}\right)^{1/2}$ a normalization constant and $P_n^{|m|}$ the associated Legendre polynomial given for $0 \leq m \leq n$ by [6]

$$P_n^m(x) := \frac{(-1)^m}{2^n n!}(1-x^2)^{m/2}\frac{\mathrm{d}^{n+m}}{\mathrm{d}x^{n+m}}(x^2-1)^n. \tag{3}$$

We consider finitely many degrees, $N \geq 0$ being the maximal degree. In particular, we have $\sum_{n=0}^{N}(2n+1) = (N+1)^2$ harmonics of degree up to N. Note that one can recover any angular pattern when $N \to \infty$. We then combine isotropic wavelets and spherical harmonics to consider filters of the form $(r,\theta,\phi) \mapsto \widehat{g}_i(r)Y_n^m(\theta,\phi)$ in Fourier domain.

SHW Invariants to Local Rotations. Fix an image f. For the scale $i \in \mathcal{I}$, the degree $n \in \{0,\dots,N\}$, and the order $m \in \{-n,\dots,n\}$, we set the texture measurement scalar $c_{i,n,m}(f) := \langle \widehat{f}, \widehat{g}_i Y_n^m \rangle$ and vector $\boldsymbol{c}_{i,n}(f) = (c_{i,n,m})_{|m|\leq n}$. Here, $\boldsymbol{c}_{i,n}(f)$ contains the spectral information for the degree n, once we have

projected the $3D$ image on the sphere using the isotropic wavelet $\widehat{g}$ at scale i. We set

$$\alpha_{i,n}(f) = \frac{\|\boldsymbol{c}_{i,n}(f)\|_2^2}{2n+1} = \frac{1}{2n+1} \sum_{-n \leq m \leq n} |c_{i,n,m}|^2, \tag{4}$$

that represents the spherical energy of the image, centered around $\mathbf{0}$, and averaged on the $(2n+1)$ components of $\boldsymbol{c}_{i,n}$, for the degree n and at scale i. The quantities (4) represent the directional information of the image of interest f at the location $\boldsymbol{x}_0 = \mathbf{0}$. The texture operators are defined by sliding the image to different center $\boldsymbol{x}_0 \in \mathbb{R}^3$. For $i \in \mathcal{I}$ and $n \leq N$, we set

$$h_{i,n}(\boldsymbol{x}_0) = \mathcal{G}_{i,n}\{f\}(\boldsymbol{x}_0) = (\alpha_{i,n}(f(\cdot - \boldsymbol{x}_0)))^{1/2}. \tag{5}$$

Here, $h_{i,n}$ is a new image (i.e. response map) that can be interpreted as follows. It extracts the energy contained in the spatial and spherical frequencies corresponding to the degree n, at scale i, for the image f centered around the location $\boldsymbol{x}_0$. Then, $\mathcal{G}_{i,n}$ is a texture operator in the sense of [3].

Proposition 1. *The texture operators $\mathcal{G}_{i,n}$ are LRI in the sense of Definition 1.*

Proof. We show (1) for $\boldsymbol{x}_0 = \mathbf{0}$. It is then easily extended by exploiting that $\mathcal{R}_{\boldsymbol{x}_0} = \mathcal{T}_{\boldsymbol{x}_0}\mathcal{R}\mathcal{T}_{-\boldsymbol{x}_0}$. The rotated spherical harmonic $\mathcal{R}Y_n^m$ is in the span of the spherical harmonics of same degree $Y_n^{m'}$; that is, $\mathcal{R}Y_n^m = \sum_{m'=-n}^{n} s_{n,\mathcal{R}}[m,m']Y_n^{m'}$. Because rotations are isometric operators, we have that

$$1 = \|Y_n^m\|_{L_2} = \|\mathcal{R}Y_n^m\|_{L_2} = \left\| \sum_{m'=-n\ldots n} s_{n,\mathcal{R}}[m,m']Y_n^{m'} \right\|_{L_2} = \|\boldsymbol{s}_{n,\mathcal{R}}[m,\cdot]\|_2. \tag{6}$$

Moreover, $\mathcal{R}Y_n^m$ and $\mathcal{R}Y_n^{m'}$ are orthogonal as soon as $m \neq m'$, implying that $\langle \boldsymbol{s}_{n,\mathcal{R}}[m,\cdot], \boldsymbol{s}_{n,\mathcal{R}}[m',\cdot]\rangle = 0$ for $m \neq m'$. In other terms, $S_{n,\mathcal{R}} = (\boldsymbol{s}_{n,\mathcal{R}}[m,m'])_{m,m'}$ is a matrix preserving the Euclidean norm.[2] We then have, using that the rotations commute with the Fourier transform and expanding $\mathcal{R}^*Y_n^m$ where $\mathcal{R}^*$ is the adjoint rotation of $\mathcal{R}$,

$$\boldsymbol{c}_{i,n}(\mathcal{R}f) = \left(\langle \mathcal{R}\widehat{f}, \widehat{g}_i Y_n^m\rangle\right)_m = \left(\langle \widehat{f}, \widehat{g}_i \mathcal{R}^* Y_n^m\rangle\right)_m = \left(S_{n,\mathcal{R}^*}\boldsymbol{c}_{i,n}(f)\right)_m. \tag{7}$$

Finally, since $S_{n,\mathcal{R}^*}$ preserves the Euclidean norm, we deduce that

$$\begin{aligned} \mathcal{G}_{i,n}\{\mathcal{R}f\}(\mathbf{0}) &= \frac{1}{2n+1}\|\boldsymbol{c}_{i,n}(\mathcal{R}f)\|_2^2 = \frac{1}{2n+1}\|S_{n,\mathcal{R}^*}\boldsymbol{c}_{i,n}(f)\|_2^2 \\ &= \frac{1}{2n+1}\|\boldsymbol{c}_{i,n}(f)\|_2^2 = \mathcal{G}_{i,n}\{f\}(\mathbf{0}). \end{aligned} \tag{8}$$

[2] The matrix $S_{n,\mathcal{R}}$ is called the steering matrix in the literature [7]. We recovered the well-known property that the steering matrix of spherical harmonics is orthogonal.

In summary, $\mathcal{G}_{i,n}$ incorporates directional information by considering different spherical frequencies defined by the degree n, and taking the norm of the vector $\boldsymbol{c}_{i,n}(f)$ allows extracting the spherical frequencies in a rotation invariant fashion. The latter can qualitatively be interpreted as the spherical Fourier modulus computed for a degree n, at a position $\boldsymbol{x}_0$ and scale i. This is the main advantage of the proposed method.

To obtain a set of scalar texture measurements, we use the average of the image $\mathcal{G}_{i,n}\{f\}$ over a VOI M as an aggregation function $\mu_{i,n}$ in the sense of [3].

3 Results

3.1 LRI and DS of Popular Radiomics Operators

We provide a qualitative comparison of popular radiomics operators[3] in terms of their DS and LRI. This comparison is presented in Table 1 and includes SHW invariants, aligned Riesz wavelets [8], separable Haar wavelets and Coiflets [4], Laplacians of Gaussians (LoG) and the three main approaches based on Gray-Level Matrices (GLM): Gray-Level Co-occurrence Matrices (GLCM) [9], Gray-Level Run-Length Matrices (GLRLM) [10], Gray-Level Size Zone Matrices (GLSZM) [11]. We refer to [3] for detailed descriptions of the texture operators. We observe that only SHW invariants and can combine LRI with DS. Aligned Riesz wavelets can approximately combine the two properties.

3.2 3D Synthetic Texture Classification

We evaluate the importance of radiomics operators combining directional sensitivity and local rotation invariance on the synthetic RFAI database [12]. The texture images resemble medical tissues observed in CT, MRI and PET, including pronounced directional components. The *Fourier* dataset includes 15 classes built from various synthetic power spectrums (see Fig. 2). There are ten $64 \times 64 \times 64$ instances per class, which are built from randomly selected phase values in the Fourier domain. Two test suites are available. The first one, referred to as *Normal* contains the initial 150 instances. A second one, referred to as *Rotate* includes 150 new instances that underwent random 3D rotations of the initial power spectrum. We use linear Support Vector Machines (SVM) in order to classify adequately the 3D textures, where the cost of errors is optimized in $[10^{-4}; 10^{6}]$. We use a Leave-One-Out (LOO) cross-validation to evaluate the average classification performance in terms of accuracy. For both *Normal* and *Rotate* test suites, only instances from the *Normal* set are used for training the models. Mirror boundary conditions and average energies of the coefficients were used for all convolutional approaches. Specific settings of the feature groups are listed in the following paragraphs.

[3] Considered popular operators are those included in radiomics libraries including pyRadiomics, TexRAD, IBEX, CERR, MAZDA, QIFE, LifeX and QuantImage.

Table 1. Qualitative comparison of popular radiomics operators.

	Directional sensitivity	Invariance to local rotations
SHW invariants	Yes, i.e. in terms of spherical frequencies of a given degree n	Yes (via Proposition 1)
Aligned Riesz wavelets [8]	Yes, i.e. in terms of higher order partial image derivatives	Approximately, via the local alignment of the filters based on the structure tensor
Haar wavelets and Coiflets [4]	Yes. The separable filters are directional and aligned with image axes	No, rotations will redistribute energies among the subbands
LoGs	No. The convolutional texture operator $g(\boldsymbol{x}) = g(\|\|\boldsymbol{x}\|\|) = g(r)$ is isotropic	Yes (via isotropy)
GLCMs [9]	No. The texture operators become isotropic when averaged over all image directions	Approximately (via discrete isotropy)
GLRLMs [10]	No. The texture operators become isotropic when averaged over all image directions	Approximately (via discrete isotropy)
GLSZMs [11]	No, the texture operator will mix elongated and circular zones	Approximately (via discrete isotropy)

SHW Invariants. One considers the feature vector $\boldsymbol{\mu} = (\mu_{i,n})_{i\in\mathcal{I},\, 0\leq n\leq N}$ from the SHW invariants. The number of features is therefore $|\mathcal{I}|\times(N+1)$, the number of scales $|\mathcal{I}|$ being chosen as 4. Concatenated degrees up to $N = 11$ are tested, resulting in feature vectors of dimensionalities from 4 to 48.

Aligned Riesz wavelets. We compute features from aligned 3D Riesz wavelets as developed in [8] and included in the QuantImage platform[4]. The Meyer isotropic wavelet is used with four scales. The variance of the Gaussian window regularizing the structure tensor is fixed to 1. Separated Riesz orders in $[0; 4]$ are tested, resulting in feature vectors of dimensionalities from 4 to 60.

[4] https://radiomics.hevs.ch, as of June 2018.

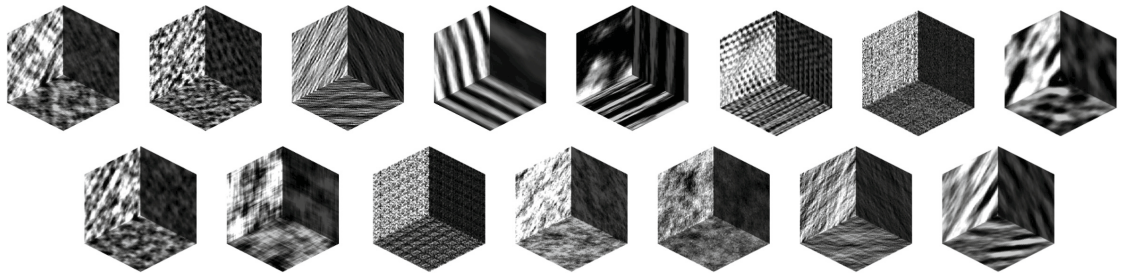

Fig. 2. The 15 classes of the *Fourier* dataset of the RFAI database [12]. Each instance is a 3D image of $64 \times 64 \times 64$ voxels.

Haar wavelets and Coiflets. 3D Separable Haar wavelets and Coiflets [4] are computed using the pyRadiomics toolbox[5] [1]. The pure Lowpass (L) filter is not used, resulting in 7 filters involving the Highpass (H) per scale (i.e. following $x-y-z$ convention: HHH, HLL, LHL, LHH, LLH, HHL, HLH). Coiflets with 1 to 5 Vanishing Moments (VM) were tested, where one VM provided best results. Concatenated scales in $[1;4]$ are tested, resulting in 7 to 28 features.
Laplacians of Gaussian. 3D LoGs are computed using pyRadiomics [1]. Concatenated variances $\sigma = 1, 1.5, \ldots, 5$ are tested, resulting in 1 to 9 features.
Gray level matrices. GLMs features including GLCMs, GLRLMs and GLSZMs are computed using pyRadiomics [1]. Default values are used. 23 GLCMs features are computed as listed in pyRadiomics' documentation. GLCMs operators with distances in $[1;5]$ are tested resulting in 23 to 115 features when concatenated. 16 GLRLMs and 16 GLSZMs features are computed using default values.

The classification accuracies with respect to the number of features are reported in Fig. 3 for *Normal* and *Rotate* test suites. SHW invariants provide excellent classification performance for a low number of features (i.e. 16 with $N = 3$). They also showed high generalization abilities with best performance on the *Rotate* dataset. The degradation of classification performance is caused by (i) the use of interpolation for implementing rotations in RFAI and (ii) important boundary effects when using four dyadic wavelet scales on $64 \times 64 \times 64$ images, which are more or less present based on a given rotation. Separable wavelets showed poor generalization on *Rotate*, which is due to their lack of LRI. LoGs, and all GLM designs achieved relatively poor performance on both datasets, which highlights the importance of DS. While LoGs showed excellent generalization on *Rotate*, all GLMs did not perform well on rotated instances. This suggests that approximate LRI via discrete isotropy is not sufficient. Overall, the results strongly support our hypothesis that combining LRI with DS is essential for optimal texture characterization.

We did not included CNNs in the comparison as the number of features largely exceeds the dimensionalities considered here and they require many more

[5] https://pyradiomics.readthedocs.io, as of June 2018.

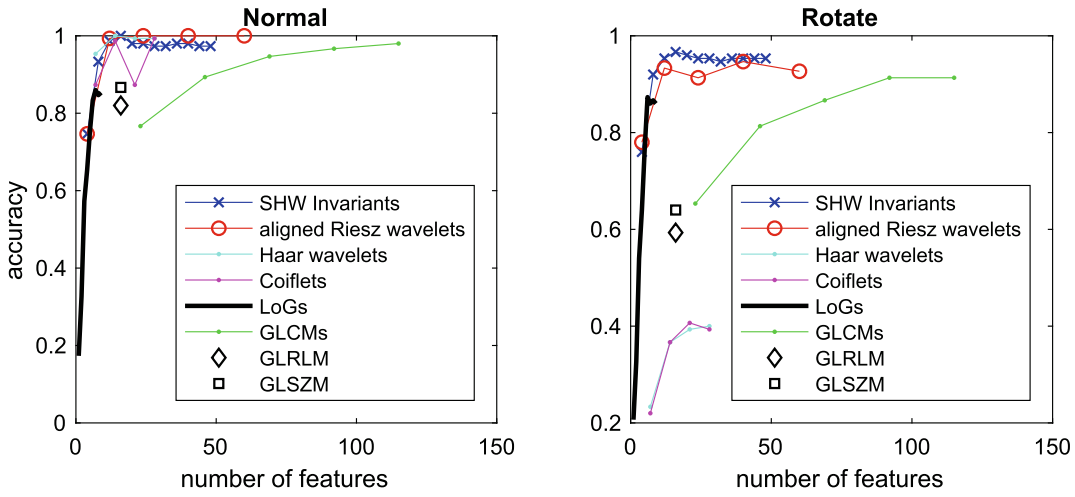

Fig. 3. Classification accuracies with respect to the number of features for both *Normal* and *Rotate* test suites. Chance level accuracy is 0.07.

instances for training. However, it is worth noting that although classical CNNs are DS but not LRI, recent designs such as group equivariant CNNs can approximate LRI [13].

4 Conclusions

We introduced the notions of LRI and DS and their importance for designing optimal radiomics image operators. The difficulty of combining the two properties was highlighted, and we subsequently proposed a simple approach based on SHW invariants to fulfill the two requirements. A qualitative comparison of popular radiomics operators revealed that none of them are able to combine LRI and DS. We demonstrated the importance of these two properties for 3D synthetic texture classification, where the proposed SHW invariants showed improved classification performance on rotated images. We believe that such operators capable of combining LRI and DS will yield optimal performance for radiomics.

Acknowledgments. This work was supported by the Swiss National Science Foundation (grants PZ00P2_154891 and 205320_179069).

References

1. van Griethuysen, J.J.M., et al.: Computational radiomics system to decode the radiographic phenotype. Cancer Res. **77**(21), e104–e107 (2017)
2. Gatenby, R.A., Grove, O., Gillies, R.J.: Quantitative imaging in cancer evolution and ecology. Radiology **269**(1), 8–14 (2013)
3. Depeursinge, A., Al-Kadi, O.S., Mitchell, J.R.: Biomedical Texture Analysis: Fundamentals. Applications and Tools. Elsevier-MICCAI Book series. Academic Press, London (2017)

4. Daubechies, I.: Ten Lectures on Wavelets, vol. 61. SIAM, Philadelphia (1992)
5. Driscoll, J.R., Healy, D.M.: Computing fourier transforms and convolutions on the 2-sphere. Adv. Appl. Math. **15**(2), 202–250 (1994)
6. Abramowitz, M., Stegun, I.: Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables, vol. 55. Courier Corporation, New York (1964)
7. Ward, J.P., Unser, M.: Harmonic singular integrals and steerable wavelets in $L_2(\mathbb{R}^d)$. Appl. Comput. Harmon. Anal. **36**(2), 183–197 (2014)
8. Dicente Cid, Y., Müller, H., Platon, A., Poletti, P.-A., Depeursinge, A.: 3-D solid texture classification using locally-oriented wavelet transforms. IEEE Trans. Image Process. **26**(4), 1899–1910 (2017)
9. Haralick, R.M.: Statistical and structural approaches to texture. Proc. IEEE **67**(5), 786–804 (1979)
10. Galloway, M.M.: Texture analysis using gray level run lengths. Comput. Graph. Image Process. **4**(2), 172–179 (1975)
11. Thibault, G., et al.: Texture indexes and gray level size zone matrix: application to cell nuclei classification. Pattern Recognit. Inf. Process. 140–145 (2009)
12. Paulhac, L., Makris, P., Ramel, J.-Y.: A solid texture database for segmentation and classification experiments. In: 4th International Conference on Computer Vision Theory and Applications, pp. 135–141 (2009)
13. Andrearczyk, V., Depeursinge, A.: Rotational 3D texture classification using group equivariant CNNs. In submitted to SPIE Medical Imaging (2018)