

CHAPTER 2

Multi-Scale and Multi-Directional Biomedical Texture Analysis

Finding the Needle in the Haystack

Adrien Depeursinge^{a,*,**}

* École Polytechnique Fédérale de Lausanne, Biomedical Imaging Group, Lausanne 1015, Switzerland.

** University of Applied Sciences Western Switzerland (HES-SO), Institute of Information Systems, Sierre 3960, Switzerland.

^a Corresponding: adrien.depeursinge@hevs.ch

Contents

1. Introduction	38
2. Notations	39
3. Multi-scale image analysis	40
3.1. Spatial versus spectral coverage of linear operators: The uncertainty principle	40
3.2. Region of interest and response map aggregation	42
4. Multi-directional image analysis	46
4.1. The Local Organization of Image Directions (LOID)	49
4.2. Directional sensitivity of texture operators	50
4.3. Locally rotation-invariant operators and moving frames representations	51
4.4. Directionally insensitive, sensitive, and moving frames representations for texture classification: a quantitative performance comparison	54
5. Discussions and conclusions	56
References	62

Abstract

This chapter clarifies the important aspects of biomedical texture analysis under the general framework introduced in Chapter 1. It was proposed that any approach can be characterized as the combination of local texture operators and regional aggregation functions. The type of scale and directional information that can or cannot be modeled by categories of texture processing methods is revealed through theoretic analyses and experimental validations. Several key aspects are found to be commonly overlooked in the literature and are highlighted. First, we demonstrate the risk of using regions of interest for aggregation that are regrouping tissue types of different natures. Second, a detailed study of the type of directional information important for biomedical texture characterization suggests that fundamental properties lie in the local organization of image directions. In addition, it was found that most approaches cannot efficiently characterize the latter, and even fewer can do it

with invariance to local rotations. We conclude by deriving novel comparison axes to evaluate the relevance of biomedical texture analysis methods in a specific medical or biological applicative context.

Keywords: Texton, moving frames, uncertainty principle, texture analysis.

1. Introduction

The diversity of existing Biomedical Texture Analysis (BTA) approaches illustrates the various properties required in different applicative contexts [14, 39]. Desired BTA properties are *e.g.*, ease of use, interpretability, low computational cost, and most importantly high discriminatory performance and specificity. The latter is strongly dependent on the nature of the texture information required for a specific task in hand. The purpose of this work is to dissect the wide range of BTA properties to provide a set of comparison dimensions between approaches. A formal definition of biomedical texture information was proposed in Section 2 of Chapter 1. The latter was found to be characterized by the type of spatial transitions and dependencies between pixel values. In particular, it was demonstrated that the scales (*i.e.*, speed of variation or frequency) and directions of the spatial transitions are fundamental properties of biomedical texture functions. In Section 3.1 of Chapter 1, a general problem formulation for biomedical texture analysis was introduced, considering that any approach can be characterized as a set of texture operators and aggregation functions. The operators allow locally isolating desired texture information in terms of spatial scales and directions of a texture image. The application of the operators over all positions of the image yields translation-equivariant feature maps containing every local responses of the latter. Scalar-valued texture measurements are obtained by aggregating feature maps over regions of interest.

In this chapter, clarifications are provided on possible design choices in terms of spatial scales and directions for operators and aggregation functions. An excellent description of the problem can be found in [48], which concerns medical image analysis in general. Our aim is to discuss the particularities of multi-scale and multi-directional analysis for biomedical texture characterization. A focus is made on linear operators. However, although the theoretic concepts presented are only valid for the latter, they should also provide intuition for designing nonlinear operators wherever possible. Multiple examples and toy problems will be provided to illustrate the concepts introduced. Among these, the uncertainty principle, a theoretic limitation of the trade-off between operator scale and locality is first recalled to provide guidelines for optimal design of operator scales. The influence of the size and shape of the region of interest

used for aggregation on texture classification and segmentation is demonstrated. The latter motivates the creation of digital tissue atlases of organs or tumors, providing powerful models of digital phenotypes. In Section 3.3 of Chapter 1 and in Chapter 7, the importance of approaches that are robust to rigid transformations (translation and rotation) was emphasized. In the second part of this chapter (Section 4), clarifications are made on directional information types that are important for biomedical texture analysis. In particular, the Local Organization of Image Directions (LOID: how directional structures intersect) are found to be fundamental. Characterizing the latter with invariance to local rotations raises several challenges. In this context, operators that are insensitive to image directions (called *circularly/spherically symmetric*) are compared to their *directional* counterparts. The destructive effect of aggregation on the ability of directional operators to characterize the LOIDs is demonstrated and motivates the use of Moving Frame (MF) texture representations. The latter consist of locally adapting a coordinate frame (*e.g.*, a set of non-collinear operators) based on an alignment criteria that is consistent¹ for all positions in the texture image. We provide evidence that MF representations allow detailed characterizations of the LOIDs with invariance to rigid transformation. A quantitative performance comparison of circularly symmetric, directional and MF texture representations for 2D texture classification is presented in Section 4.4. Finally, most important aspects of operator and aggregation function design are summarized under the form of a checklist matrix in Section 5.

2. Notations

Additional notations are introduced based on the notations initially defined in Section 2.2 of Chapter 1. To further analyze the relationship between scales and directions, let us consider the definition of texture functions in 2D polar coordinates as

$$\begin{array}{ccc}
 f(r, \theta) & \xleftrightarrow{\mathcal{F}} & \hat{f}(\rho, \vartheta) \\
 r \in \mathbb{R}^+, \theta \in [0, 2\pi) & & \rho \in \mathbb{R}^+, \vartheta \in [0, 2\pi) \\
 \mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = r \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix} & & \boldsymbol{\omega} = \begin{pmatrix} \omega_1 \\ \omega_2 \end{pmatrix} = \rho \begin{pmatrix} \cos \vartheta \\ \sin \vartheta \end{pmatrix}
 \end{array}$$

¹A simple and reliable alignment criteria is to orient all operators at each image position with the direction that maximizes the local image gradient.

and in 3D spherical coordinates as

$$\begin{array}{ccc}
 f(r, \theta, \phi) & \xleftrightarrow{\mathcal{F}} & \hat{f}(\rho, \vartheta, \varphi) \\
 r \in \mathbb{R}^+, \theta \in [0, \pi], \phi \in [0, 2\pi) & & \rho \in \mathbb{R}^+, \vartheta \in [0, \pi], \varphi \in [0, 2\pi) \\
 \mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = r \begin{pmatrix} \sin \theta \cos \phi \\ \sin \theta \sin \phi \\ \cos \theta \end{pmatrix} & & \boldsymbol{\omega} = \begin{pmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{pmatrix} = \rho \begin{pmatrix} \sin \vartheta \cos \varphi \\ \sin \vartheta \sin \varphi \\ \cos \vartheta \end{pmatrix}
 \end{array}$$

Polar and spherical representations allow separating the *angular* part (*i.e.*, image directions) from the *radial* part (*i.e.*, spatial frequencies related to image scales) [55]. For simplifying the notations, all coordinate domains are considered continuous in this chapter. Their discretized versions can be obtained following the notions introduced in Section 2.2 of Chapter 1.

3. Multi-scale image analysis

The need for multi-scale image analysis is motivated throughout various chapters of this book (*e.g.*, Section 3 of Chapter 1) as well as Chapters 4, 5 and 7. In this section, we will define more precisely the important aspects of multi-scale texture operator and aggregation function design. In particular, the discussion on which scales are optimal for biomedical texture measurements hinges on two important facets:

- How to optimally define the spatial support(s) $\mathbf{G}_n = G_{1,n} \times \cdots \times G_{D,n}$ and the radial responses of the operator(s) \mathcal{G}_n ?
- What is the best size and shape of the region of interest \mathbf{M} for aggregation ?

These two aspects are illustrated in Fig. 2.1 and detailed in the following subsections 3.1 and 3.2.

3.1. Spatial versus spectral coverage of linear operators: The uncertainty principle

Linear texture operators are expected to be *band-pass* functions, which means that the covered spatial frequencies are in a range defined by $\rho_{\min} > 0$ and $\rho_{\max} < \pi$ in the Fourier domain. Band-pass operators are commonly used as operators for texture analysis because their geometric behavior, isotropy or directionality can be well controlled. Because they do not include the zero frequency $\rho = \|\boldsymbol{\omega}\| = 0$, they are only sensitive to transitions between pixel values (*i.e.*, texture) and not to the average regional intensity. In between 0 and π , the *uncertainty principle* allows defining rules to design texture operators with optimal spectrum coverage $[\rho_{\min}, \rho_{\max}]$ (see Eq. 2.1). Ideal texture operators would be accurately localized both in spatial and Fourier domains. On the one hand, well localized operators in the spatial domain allow identifying precise local texture properties without including surrounding image structures.

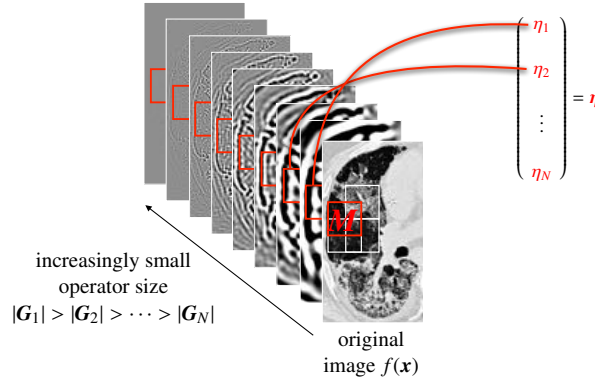


Figure 2.1: Aspects of multi-scale texture operator and aggregation function design [17]. How to optimally define the sizes $|\mathbf{G}_1| > |\mathbf{G}_2| > \dots > |\mathbf{G}_N|$ and the radial responses of a collection of operators \mathcal{G}_n ? What is the best position, size and shape of the region of interest \mathbf{M} for aggregation?

On the other hand, optimally localized operators in the Fourier domain can precisely characterize narrow frequency bands without mixing with other neighboring spectral components. Unfortunately, having both properties together is not possible and subject to a theoretic limitation called the *uncertainty principle* [40]. Intuitively, the latter can be understood as follows: it is impossible to measure rich texture information from gray-level transitions between a few pixels only. Likewise, measuring all transitions from a large number of pixels yields detailed texture information, but requires large image neighborhoods. More precisely, the relationship between the spatial support $\mathbf{G} = G_1 \times \dots \times G_D$ and the spectral support $\mathbf{\Gamma} = \Gamma_1 \times \dots \times \Gamma_D$ of a linear texture operator is

$$\prod_{d=1}^D G_d^2 \Gamma_d^2 \geq \frac{1}{4^D}. \quad (2.1)$$

This trade-off is illustrated in Fig. 2.2 for 2D circularly symmetric Gaussian operators $g_n(\mathbf{x})$. Therefore, an operator with an accurate spatial localization (*i.e.*, narrow support) yields poor spectrum estimates. This has a direct implication in practice, and can be critical when the texture processes are multi-spectral and highly non-stationary. The spatial support of the operator needs to be large enough to accurately characterize the frequency components of the intricate texture process, and can potentially be larger than the studied texture region. This is illustrated in Fig. 2.3 for the characterization of local ground glass and reticular regions in lung CT. In addition, controlling the profile of the operators (*i.e.*, the decay slope at their boundaries) is important to avoid extensive *ringing effects* in their dual representation resulting in poorly localized analysis




Figure 2.2: Uncertainty principle: the function of a linear operator $g_n(\mathbf{x})$ (2D circularly symmetric low-pass Gaussian in this example) cannot be well localized both in spatial and Fourier domains. The theoretic limit observes $G_1^2 \Gamma_1^2 G_2^2 \Gamma_2^2 \geq \frac{1}{16}$.

(see Fig. 2.4) [15]. Finding the optimal trade-off between the accurate definitions of operator supports in space and in Fourier requires identifying spatial frequencies that are important for the texture segregation task in hand. The lowest discriminative frequency will determine the smallest operator size needed to differentiate between the various texture classes (see Fig. 2.5). The latter is not straightforward in most cases and machine learning can be used to determine discriminative scales [30].

3.2. Region of interest and response map aggregation

Another critical aspect of scale definition in biomedical texture analysis concerns the design of the ROI for aggregating the operators’ response maps (see Fig. 2.1). The fundamental underlying question is: how large must be the ROI \mathbf{M} ? Addressing this issue requires considering once more the spectral complexity and spatial stationarity of the considered texture processes. On the one hand, \mathbf{M} should be large enough to capture the discriminative statistics of the operators’ responses. On the other hand, using large \mathbf{M} covering several contiguous interleaving non-stationary processes will mix the statistics of the latter and result in meaningless texture measures, even when using appropriate texture operators.

Two examples are developed to illustrate the impact of the size of \mathbf{M} on texture classification and segmentation (see Figures 2.6 and 2.7). For both examples, simple circularly symmetric band-pass and multi-scale operators are used. They are based on two consecutive dyadic iterations of Simoncelli wavelet frames [41] $g_n(\mathbf{x})$. The 2D




Figure 2.3: Challenges of operator design for complex and non-stationary biomedical texture processes. Top row (peripheral ground glass opacities in chest CT): large influence of proximal objects when the support of operators is larger than the region of interest. The lung boundary has an increasingly important impact in the peripheral region, which can be observed on the response maps $h_n(x)$ of increasingly large LoG linear operators with functions $g_n(x)$ [15]. Bottom row (reticular and normal lung parenchyma in CT): on the left image, a small-sized Gaussian filter g_1 ($\sigma = 3.2\text{mm}$) is precisely located in the reticular pattern but yields a poor characterization of the spectral content inside its support. Conversely, the right image shows a large Gaussian filter g_2 ($\sigma = 38.4\text{mm}$) allowing an accurate estimation of spatial frequencies, but its spatial support encroaches upon normal parenchyma and mixes properties from the two distinct texture classes.




Figure 2.4: Importance of the decay of linear operators on spatial versus spectral support (1D). Smooth operators avoid ringing effects in the dual representation.




Figure 2.5: A texture operator characterizing spatial frequencies along x_1 in $[\omega_{1,\min}, \pi)$ is enough to discriminate f_1 from f_2 . Finding this lower bound allow defining texture operators with narrow spatial supports that are optimally localized.

version is defined in Fourier in polar coordinates (ρ, ϑ) as

$$\begin{aligned} \hat{g}_1(\rho) &= \begin{cases} \cos\left(\frac{\pi}{2} \log_2\left(\frac{2\rho}{\pi}\right)\right) & \text{for } \frac{\pi}{4} < \rho \leq \pi, \\ 0 & \text{otherwise.} \end{cases} \\ \hat{g}_2(\rho) &= \begin{cases} \cos\left(\frac{\pi}{2} \log_2\left(\frac{4\rho}{\pi}\right)\right) & \text{for } \frac{\pi}{8} < \rho \leq \frac{\pi}{2}, \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \quad (2.2)$$

Because it is circularly symmetric, this operator depends on the radial coordinate ρ only and is qualitatively similar to 2D LoGs. It is applied to a texture function f with convolution (*i.e.*, equivalent to a multiplication in the Fourier domain) as $\hat{h}_n(\rho, \vartheta) = \hat{g}_n(\rho) \cdot \hat{f}(\rho, \vartheta)$.

A first example is detailed in Fig. 2.6 and demonstrates the negative impact of averaging feature maps over large ROIs including texture functions from distinct spatial processes. This is a critical issue when texture analysis is used to characterize the structural properties of tumors because requires defining smaller ROIs based on distinct tumor habitats [25] (*e.g.*, ground glass versus solid tumor components in lung adenocarcinoma [19]). When the component textures are too expensive to delineate (*e.g.*, tumor habitats in 3D imaging) or when they are not known in advance, unsupervised texture segmentation approaches can be used to reveal the diversity of patterns contained in a given ROI. Examples of such methods are superpixels [1], graph

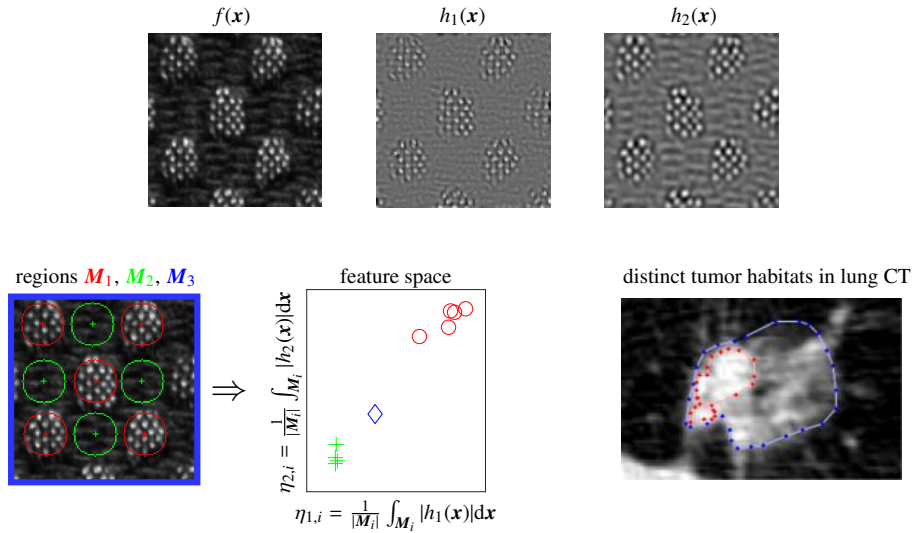


Figure 2.6: Influence of the size and localization of the region of interest M for aggregating the feature maps using the average. Bottom center: each region M_i is represented by a point (e.g., “+”) in a feature space spanned by feature averages $(\eta_{1,i}, \eta_{2,i})$. Whereas red and green regions are well separated and regrouped in the feature space, averaging the feature maps over the entire image (blue region) yields texture features that do not correspond to anything visually (see the blue diamond in the feature space). Likewise, averaging texture properties over entire tumor regions including distinct habitats will provide quantitative texture measures that do not correspond to anything biologically [25, 10].

cuts [6], or the Pott’s model [47]. An advantage of the Pott’s model is its ability to handle multiple feature maps for segmenting the subregions and therefore can easily run on the registered outputs of several operators. The aggregation function has itself a strong influence on the specificity of the texture measures. In [10], Cirujeda *et al.* showed that measurements based on the covariances of the operator’s responses provided a better characterization of texture properties in multi-component tumors when compared to the average, thanks to their ability to quantify the point-wise co-activation of the operators.

In a second example, the influence of the size of circular patch-based ROIs on supervised texture segmentation is investigated (see Fig. 2.7). Linear Support Vector Machines (SVM) are trained from overlapping ROIs extracted from two unrotated instances of classes $f_1(x)$: “canvas002” and $f_2(x)$: “canvas003” of the Outex database [37]. The responses of circularly symmetric Simoncelli wavelet frames (see Eq. 2.2) are averaged over the circular ROIs to provide texture measures $\eta = (\eta_1, \eta_2)$.

Every block is represented in the two-dimensional feature space yielded by $\text{span}(\boldsymbol{\eta})$. In the latter, SVMs learn a separating hyperplane $\boldsymbol{w}_{\text{SVM}} + b$, which is further used to classify overlapping patches from a test image composed of rotated instances of f_1 and f_2 . The corresponding decision values of the test patches (*i.e.*, $\langle \boldsymbol{\eta}, \boldsymbol{w}_{\text{SVM}} \rangle + b$) as well as the predicted local labels (*i.e.*, $\text{sgn}(\langle \boldsymbol{\eta}, \boldsymbol{w}_{\text{SVM}} \rangle + b)$) are shown for three different patch radii. The evolution of the segmentation error with radii varying in $[0, 128]$ is shown in Fig. 2.7 top right, highlighting the importance of the size of the aggregation region \boldsymbol{M} . When \boldsymbol{M} is too small (radius = 8), the local average of the feature maps is poorly estimated, which yields noisy feature estimates $\boldsymbol{\eta}$ (error = 0.24). At the other extreme, very large regions \boldsymbol{M} (radius = 128) yield accurate estimates of the features, but are not well localized spatially. This results in important errors at the boundaries between f_1 and f_2 (error = 0.23). In between these two extremes, finding adequate sizes seems not critical and allows satisfactory segmentation results with a minimum error of 0.045 for a radius of 90. However, a radius of 30 allows obtaining an excellent trade-off between locality and average estimation (error = 0.051). To summarize, a simple rule of thumb to observe is to use ROIs that are no larger than enough to accurately estimate discriminative statistics of the operators’ responses over stationary areas defined in terms of human perception or tissue biology.

In most cases, it is not realistic to assume that the texture properties are homogeneous (*i.e.*, stationary) over entire organs or tumors, which was mostly overlooked in the literature. An interesting approach is to divide organs into subregions for which it is reasonable to consider that the responses of texture operators are stationary in the relaxed sense. This provides the exciting opportunity to construct tissue atlases from texture information in biomedical images. They can be used to create disease-specific *digital phenotypes*, which already showed to constitute powerful models for predicting disease diagnosis, treatment response and/or patient prognosis in the context of interstitial lung diseases [12, 20] (see Fig. 2.8) and cancer [19, 50].

4. Multi-directional image analysis

As already introduced in Section 2.1 of Chapter 1, the notion of direction in texture is fundamental and complementary to the notion of scale (see Fig. 1.2 of Chapter 1). An important question is then: which image directions are important for deriving texture measurements and allowing adequate texture segregation? Fig. 2.9 shows textures and their Fourier modulus. For one of them, $f_1(\boldsymbol{x})$, directionality seems not obvious because the texture contains little structure and is highly stochastic. For the other, $f_2(\boldsymbol{x})$, dominant directions are clearly visible, creating oriented grid patterns and corresponding peaks in the Fourier domain. Moreover, it appears that texture directionality is defined for a particular position and scale (*i.e.*, spatial frequency). These aspects of texture directionality are developed in the next subsections 4.1, 4.2 and 4.3.

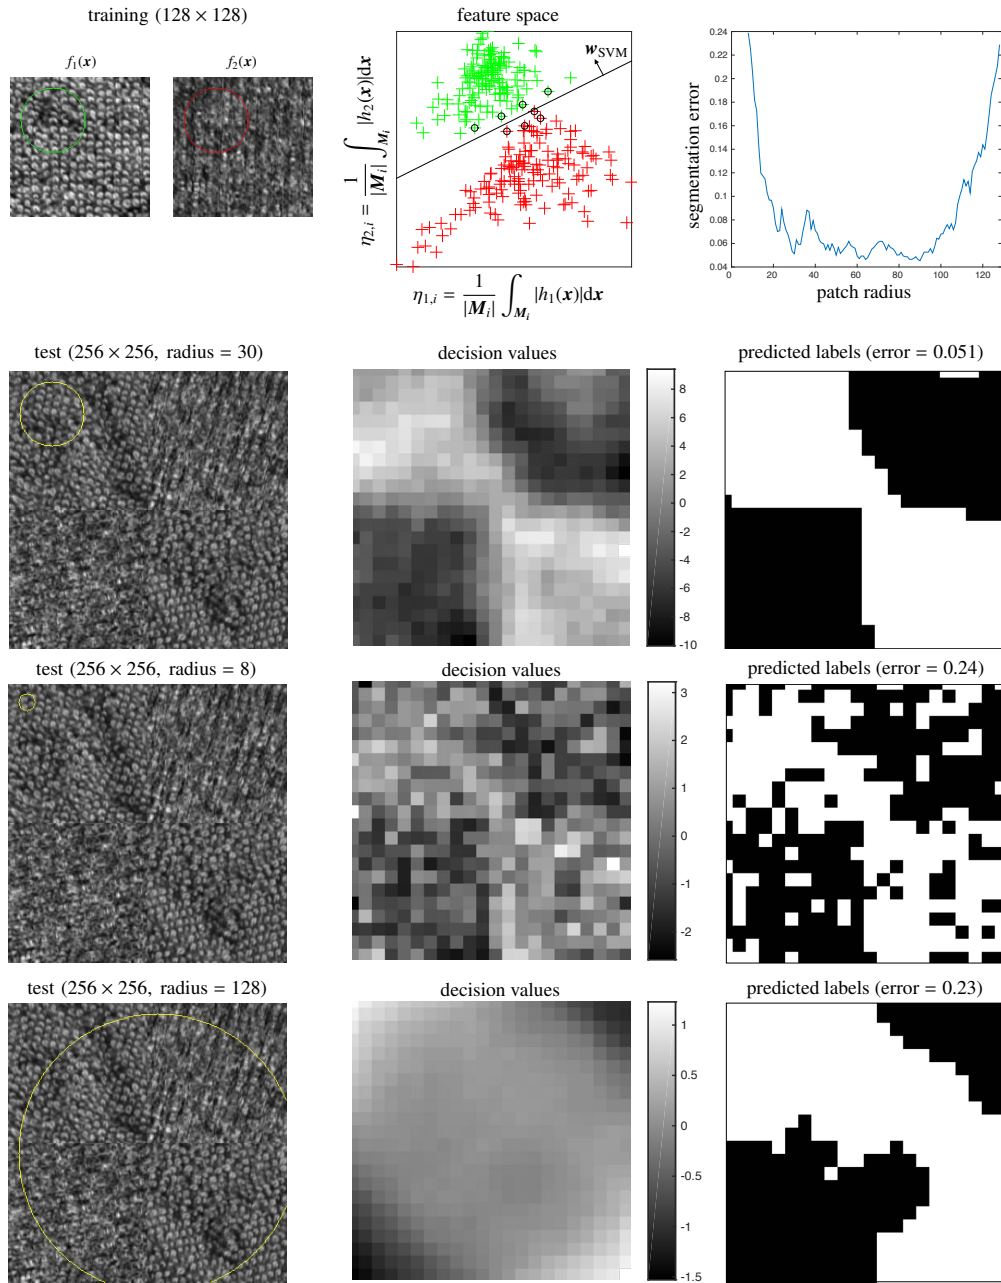


Figure 2.7: Influence of the size of circular patch ROIs on supervised texture segmentation.




Figure 2.8: Digital phenotypes for interstitial lung diseases. Left: due to the presence of thicker bronchovascular structures in the region M_a that is close to the mediastinum, the texture properties of normal and altered parenchymal tissue cannot be considered similar as the ones in the peripheral region M_b . Right: therefore, it is relevant to divide lungs into regions for which it is reasonable to consider that texture properties are homogeneous and create tissue atlases to derive digital phenotypes for interstitial lung diseases [12, 20].




Figure 2.9: Examples of texture functions with weak (*i.e.*, $f_1(x)$) versus strong (*i.e.*, $f_2(x)$) directionality. The modulus of the Fourier representation of $f_2(x)$ shows that there are clear directional spatial frequencies, standing out as bright spots in $|\hat{f}_2(\omega)|$. This also demonstrates that image directionality is defined for a particular scale (*i.e.*, spatial frequency).




Figure 2.10: Importance of the Local Organization of Image Directions (LOID) in natural and biomedical textures (*i.e.*, how directional structures intersect). a. Photograph of creased paper. b. Photomicrograph of hypertrophic cardiomyopathy [26]. c. Chest CT angiography. d. Photograph of meandroid coral. e. Fluorescence microscopy cross-sectional photograph of the tibialis anterior muscle of a mouse [7]. f. Honeycombing fibrosis in lung CT [18].

4.1. The Local Organization of Image Directions (LOID)

Thinking even further, it appears that most biomedical and natural textures have clear directional structures or primitives, but the latter are not necessarily consistent over large regions M . Most often the opposite happens where directional structures are defined locally (see Fig. 2.10). More precisely, an important aspect of directionality in natural and biomedical textures is the Local Organization of Image Directions (LOID), *i.e.*, how directional structures intersect (see Fig. 2.11). The LOIDs were already mentioned in the literature as being central in preattentive texture segregation [4] as well as computerized texture analysis [3, 16] (see Section 2.4 of Chapter 1). They relate to the primitives or textons of biomedical texture (*i.e.*, the essential “stitches” of the tissue), which often have random local orientations. The various forms of these tissue stitches are even richer in 3D, where the potential complexity of the primitives follows a cubic growth. The number of possible discrete image directions grows as $(2r + 1)^3 - 1$ in 3D versus $(2r + 1)^2 - 1$ in 2D [14]. Therefore, defining texture operators that are able to characterize the LOIDs in a locally rotation-invariant fashion are required to accurately analyze biomedical texture. Advanced methods to meet these challenging requirements are further developed in Section 4.3.




Figure 2.11: Importance of the LOIDs in preattentive texture segregation [4] (see Section 2.4 of Chapter 1). The LOIDs can be distinguished by counting the number of endpoints of the primitives (top row).

4.2. Directional sensitivity of texture operators

It is convenient to consider two distinct categories of operators in terms of directional characterization: directionally sensitive versus insensitive. In the particular case of linear operators, directionally insensitive operators are called *circularly/spherically symmetric* operators and their functions do not depend on the angular coordinate(s):

$$g_n(r) \xleftrightarrow{\mathcal{F}} \hat{g}_n(\rho). \quad (2.3)$$

Examples of such operators in 2D and 3D are Gaussian filters, LoGs and circularly symmetric wavelets [41, 53] (see Fig. 2.12). Examples of non-linear directionally insensitive operators are max or median filters. Directional operators constitute a vast category where operator functions depend on all polar/spherical coordinates. They include Fourier basis functions, circular and spherical harmonics [51, 56], directional filters and wavelets (*e.g.*, Gabor [5], Riesz [52], Simoncelli’s steerable pyramid [45], curvelets [8]), Histogram of Oriented Gradients, HOG [34, 35, 32], GLCMs [29], LBPs [38], GLRLMs [23, 57], CNNs [43], DL [24, 36], and others.

By construction, directionally insensitive operators are locally rotation-invariant, but insensitive to image directions. Texture measures obtained from this category of operators are therefore invariant to local rotations transformations. However, they can hardly differentiate between **+**-shaped, **L**-shaped or blob-shaped texture primitives (see Fig. 2.13 bottom left). Therefore, they cannot characterize the LOIDs and can only be used to distinguish between biomedical tissue types with manifest differences in image scales. Directional counterparts are sensitive to image directions, but may not be locally rotation-invariant even in an approximate sense. They are able to identify the LOIDs only when they have all the same orientation, which is very unlikely in biomedical textures (see Fig. 2.10). Moreover, the characterization of the LOIDs can be challenging even when the latter are all aligned to each other. In fact, the

aggregation function plays itself an important role when unidirectional operators² are jointly used to characterize the LOIDs [16]. When separately integrated, the responses of unidirectional individual operators are not local anymore and their joint responses become only sensitive to the global amount of image directions in the region \mathbf{M} . For instance, the joint responses of image gradients

$$\mathcal{G}_d\{f\}(\mathbf{x}_0) = \frac{\partial f}{\partial x_d}(\mathbf{x}_0), \quad d = 1, 2, \quad (2.4)$$

are not able to discriminate between the two textures classes $f_1(\mathbf{x})$ and $f_2(\mathbf{x})$ shown in Figure 2.13 when integrated over the full image domain \mathbf{M} . This loss of information is detailed by Sifre *et al.* in terms of separable group invariants [44]. When integrated separately, the responses of unidirectional operators become invariant to a larger family of roto-translations where different orientations are translated by different values. For instance, it can be observed in Fig. 2.13 that f_2 can be obtained from f_1 by vertically translating horizontal bars only and horizontally translating vertical bars only.

Further refinements are required to allow for a true locally rotation-invariant characterization of the LOIDs. Several approaches were proposed to increase the local rotation-invariance of directional operators. GLCMs and GLRLMs are made approximately insensitive to directions either by averaging feature measures or by summing the counts over all directions of the operators [49] (*e.g.*, co-occurrences along x_1 or x_2 are mixed). Likewise, rotation-based data augmentation in CNNs and DL improves invariance to local rotations [27]. Unfortunately, these processes reduce the ability of operators to characterize the LOIDs by making them insensitive to image directions. More advanced approaches were developed to allow enhanced locally rotation-invariant characterization of the LOIDs and are described in Section 4.3.

4.3. Locally rotation-invariant operators and moving frames representations

It was observed in Section 4.2 that simple texture operators are either locally rotation-invariant and directionally insensitive or able to characterize the LOIDs (directionally sensitive), but regrouping both properties is not straightforward. A variety of texture analysis approaches have been proposed to tackle this challenge. Those include the Maximum Response 8 (MR8) filterbank [54], rotation-invariant LBPs [2, 38] and extensions [28, 33], discrete and continuous HOGs (*e.g.*, used in the Scale-Invariant Feature Transform, SIFT and in the Rotation-Invariant Feature Transform, RIFT) [35, 32, 31], as well as oriented [58] and steerable [13, 16, 21] filters and wavelets that were also included in recent CNNs [27, 44]. All the aforementioned approaches rely

²Unidirectional operators are “seeing” only one direction.

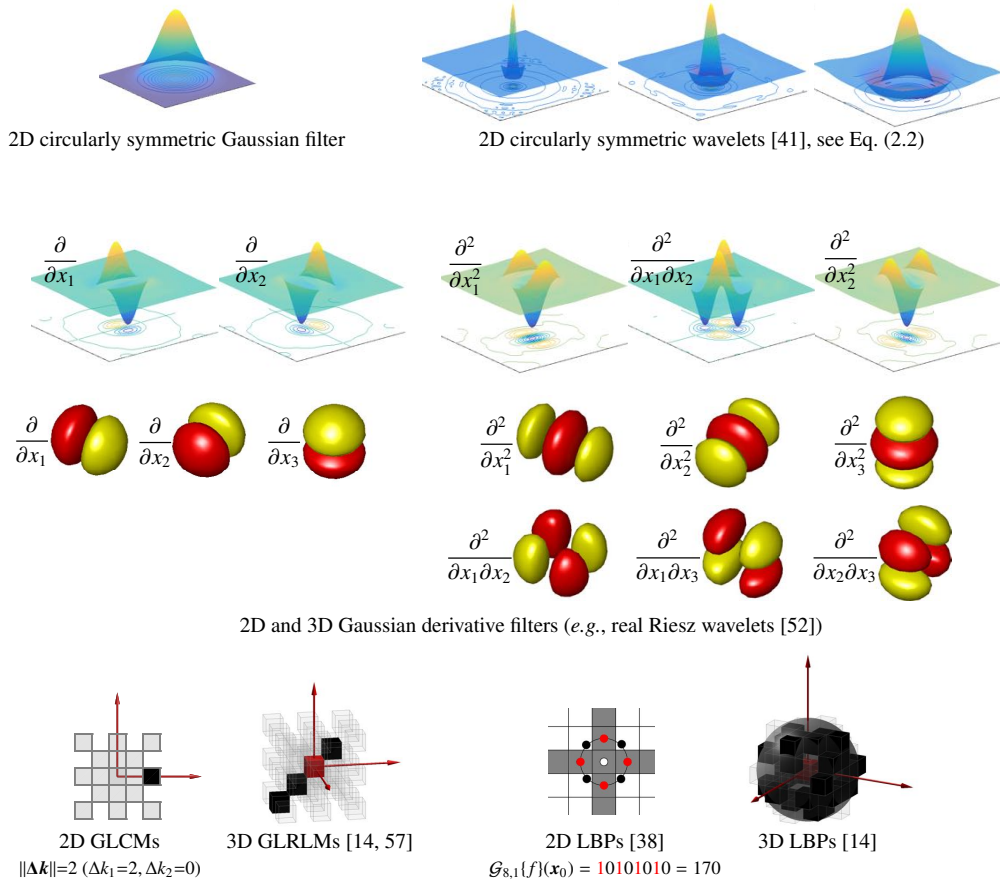


Figure 2.12: Directionally sensitive versus insensitive texture operators. Top row: linear circularly symmetric operators that are not sensitive to image directions and therefore locally rotation-invariant. Middle rows: linear directionally sensitive operators. Bottom row: non-linear directionally sensitive operators. Directional operators are sensitive to image directions but not locally rotation-invariant by construction.

on the same strategy: using operators \mathcal{G}_n that are (approximately) locally rotation-equivariant over their own support \mathbf{G}_n and then to align the operators to achieve local rotation-invariance.

Locally rotation-invariant characterization of the LOIDs can be efficiently carried out using Moving Frames (MF) representations [3, 16]. The key idea of MFs is to locally adapt a coordinate frame directly to image contours instead of using fixed extrinsic coordinates for all image locations (see Figure 2.14). The two necessary and sufficient requirements to define MFs are (i) using a set of $N \geq D$ non-collinear tex-

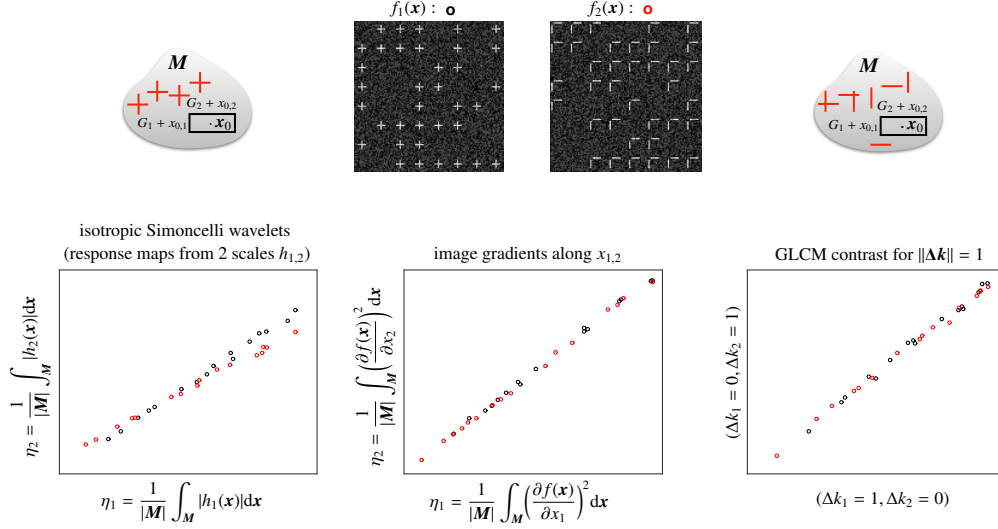


Figure 2.13: $f_1(\mathbf{x})$ and $f_2(\mathbf{x})$ only differ in terms of the LOIDs (*i.e.*, +-shaped versus L-shaped). Bottom row: one circle in each feature representation corresponds to one realization (*i.e.*, full image) of f_i , where feature maps are averaged over the entire image. Bottom left: feature vectors η_i obtained from the responses of circularly symmetric operators (two consecutive dyadic scales of Simoncelli wavelets, see Eq. 2.2) provide poor distinction between the two classes. Bottom center and right: even when the LOIDs are all aligned to each other, the joint responses of directional operators (*e.g.*, image gradients along x_d , co-occurrences along x_d) can hardly discriminate between f_1 and f_2 when integrated over the full image domain \mathbf{M} .

ture operators and, (ii) having a consistent criteria to define the local orientation of the frame bundle (*e.g.*, using the tangent as the first unit vector of the frame). In 2D, orthonormal MFs are defined as

$$\mathbf{e}_{1,x_0} = \cos \theta_{x_0} \cdot \mathbf{e}_1 + \sin \theta_{x_0} \cdot \mathbf{e}_2, \quad (2.5)$$

$$\mathbf{e}_{2,x_0} = \cos(\theta_{x_0} + \pi/2) \cdot \mathbf{e}_1 + \sin(\theta_{x_0} + \pi/2) \cdot \mathbf{e}_2, \quad (2.6)$$

where $\{\mathbf{e}_1, \mathbf{e}_2\}$ is the canonical basis for \mathbb{R}^2 , $\{\mathbf{e}_{1,x_0}, \mathbf{e}_{2,x_0}\}$ is a local moving frame bundle for the position \mathbf{x}_0 , and θ_{x_0} its orientation relatively to $\{\mathbf{e}_1, \mathbf{e}_2\}$. Image representations obtained from MFs are robust to rigid transformations [22] (proof in [16]). They are invariant to local rotations and equivariant to translations. Moreover, deriving the local orientation of the frame tends to preserve the joint information between positions and orientations even when the operators are integrated (*e.g.*, averaged) over an image domain \mathbf{M} . Suitable local orientation measures for defining a consistent MF alignment criterion θ_{x_0} are *e.g.*, simple pixel differences of a Gaussian-smoothed response




Figure 2.14: Construction of MFs in a histopathological image of dense connective tissue (left). Right: MFs are based on local directions θ_{x_0} maximizing the gradient magnitude (center).

map [35] (see Eq. 3.7 of Chapter 3), or the Gaussian-smoothed Hessian-based structure tensor, which can be interpreted as a localized covariance matrix of the gradient [42] (2D [46], 3D [9]). The construction of 2D MFs based on directions maximizing the gradient is illustrated in Fig. 2.14. It can be observed in Fig. 2.15 that when compared to using global coordinates (see Fig. 2.13), texture measures obtained from identical operators (*e.g.*, image gradients) but expressed in MFs can provide very detailed characterizations of the LOIDs.

It is important to note that image representations based on locally rotation-invariant operators \mathcal{G}_n are not preserving image layouts (*i.e.*, large-scale organization of image structure) larger than their spatial supports $G_{1,n} \times \dots \times G_{D,n}$. They cannot be used alone to characterize natural images with well-defined global layouts such as in ImageNet [11]. They are best suited for discriminating textured patterns with well-pronounced local directional structures, which is the case for most biomedical tissue architectures. Tuning the support of the operators and/or using hybrid approaches based on both local and global image transforms can be required to achieve optimal texture discrimination.

4.4. Directionally insensitive, sensitive, and moving frames representations for texture classification: a quantitative performance comparison

In order to reveal the differences in terms of texture discrimination performance between directionally insensitive, sensitive, and MF representations, a quantitative comparison is proposed in this section.

The Outex database of 2D natural textures with highly controlled imaging conditions [37] is chosen to compare between the various representations. The Outex_TC_10 test suite is used, because it has similar properties to biomedical images: it contains no significant changes in terms of image scale and illumination. The texture classes contain strongly directional patterns. Moreover, the validation scheme allows train-




Figure 2.15: Using gradient-based MF representations to discriminate texture classes that only differ in terms of their LOIDs. Because, the MF bundle $\{\mathbf{e}_{1,x_0}, \mathbf{e}_{2,x_0}\}$ is locally aligned with the direction θ_{x_0} maximizing the gradient at each position \mathbf{x}_0 , the energies of gradients along \mathbf{e}_{2,x_0} are null except at the center of $+$ -shaped primitives. The MF representation yields a linearly separable feature representation of f_1 and f_2 , while the same unidirectional texture operator pair (*i.e.*, orthogonal image gradients) used in global coordinates could not distinguish between the two (see Fig. 2.13). Bottom row: the flip side of the coin is that MF representation cannot differentiate textures that only differ in terms of the orientations of their primitives (*e.g.*, f_1 versus \tilde{f}_1).

ing on unrotated images only, but the testing set contains rotated instances only. This allows evaluating two important properties of the texture representations: rotation-invariance and their ability to characterize directional patterns. The classes are very pure though, where little intra-class variations are present, which differs from most biomedical texture analysis problems. Outex_TC_10 contains 24 classes, which are depicted in Fig. 2.16. It has a total of 4320 ($24 \cdot 20 \cdot 9$) non-overlapping 128×128 image instances. The training set consists of the 480 ($24 \cdot 20$) unrotated images and the remaining 3840 ($24 \cdot 20 \cdot 8$) images from 8 different orientations are constituting the test set.

Riesz wavelet frames are used as texture operators, because variations in their design allow implementing all three representation types that we want to compare [13].

The latter are detailed in Section 2.2 of Chapter 3. Qualitatively, Riesz wavelet frames correspond to multi-scale directional image derivatives and evaluate not only the magnitude, but also the type of transitions between image pixels (*i.e.*, derivative order such as the gradient, Hessian). Moreover, Riesz wavelets are steerable, which means that it is relatively inexpensive to locally align every texture operators in order to obtain rich MF representations. Four iterations of Simoncelli’s circularly symmetric wavelets (see Eq. 2.2) are used to define the spatial supports of Riesz-based image derivatives. In order to obtain MF representations, the angle θ_{x_0} maximizing the response of the first element of the filterbank $\mathcal{G}_{\sigma,L,0}\{f\}(\mathbf{x}_0)$ is used to define the MF alignment criteria at each position \mathbf{x}_0 (see *e.g.*, Eq. 3.14 of Chapter 3). A collection of texture measurements $\boldsymbol{\eta}$ is obtained by averaging the energies of the responses of each operator over the full 128×128 support \mathbf{M} of the texture instances. Simple one-versus-all SVM models using Gaussian kernels are used to learn decision boundaries in the space spanned by $\boldsymbol{\eta}$ for texture classification.

The performance comparison for various orders L of the Riesz transform is shown in Fig. 2.16. The performance of directionally sensitive unaligned Riesz wavelets appears to be clearly inferior to both directionally insensitive and MF representations. This can be explained by their lack of rotation invariance, where rotations of the texture instances swaps the responses between filters and results in noisy and diffuse class representations in the feature space. Directionally insensitive representations (*i.e.*, using circularly symmetric Simoncelli wavelets only where Riesz order $L = 0$) benefit from their invariance to local rotations and achieve an honorable classification accuracy of 87.5% with a very simple approach and low computational complexity. Best results are obtained by MF representations with an accuracy of 97.42% for $L = 4$, which highlights the importance of the LOIDs in texture analysis. However, when compared to directionally insensitive operators, MFs involve a much higher computational cost to locally estimate θ_{x_0} and align the Riesz frame accordingly. Therefore, the success of MF representations will depend on how important the LOIDs are to characterize the considered biomedical texture classes, which will set the threshold to invest the extra computational cost required. Other techniques relating to tissue underlying physiology and allowing rotational invariance are discussed in Chapters 3, 5 and 7.

5. Discussions and conclusions

This chapter explained the essential theoretic foundations and practical consequences of texture operator and aggregation function design in terms of image scales and orientations. A set of comparison dimensions was introduced, which can be used to evaluate the relevance of a particular BTA approach or design for a given biomedical application. The most important aspects are recalled in a checklist matrix (see Fig. 2.17). The expression of operators and texture functions in polar and spherical coordinates

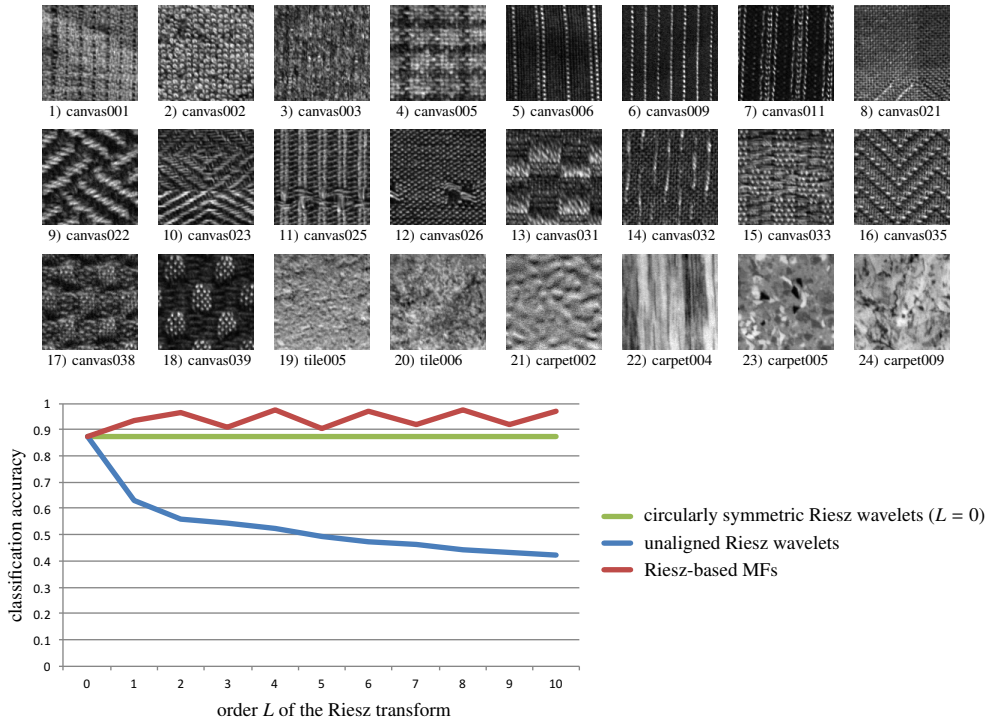


Figure 2.16: Quantitative comparison of directionally insensitive, sensitive, and MF representations for texture classification. Whereas MF representations achieve best performance, directionally insensitive texture operators allow close classification accuracies with a much lower computational complexity. Finally, unaligned directionally sensitive Riesz filters perform poorly because they lack rotation-invariance.

allowed to clearly separate scale from directional considerations.

Section 3 detailed the importance and consequences of appropriate choices of scale for operators and region of interest (see Sections 3.1 and 3.2, respectively). The uncertainty principle, a fundamental theoretic limitation was recalled to make the relation between operator locality in space and frequency explicit (see Eq. 2.1 and Fig. 2.2). As a rule of thumb, the operators should be kept as small as possible in the spatial domain to allow local texture analysis. However, they should be sufficiently large to allow for an adequate and accurate characterization of texture frequency components (see Figures 2.3 bottom row and 2.5). The importance of using operators with smooth profiles was highlighted to obtain an optimal trade-off between localization in the spatial and Fourier domains [48] (see Fig. 2.4). In addition, operator smoothness allows limiting the effect of proximal objects surrounding the region of interest for texture

analysis, which was illustrated in Fig. 2.3 (top row). The influence of the design and shape of the aggregation region was detailed in Section 3.2. In particular, the hazard of using large ROIs encompassing multiple non-stationary texture processes for aggregation with integrative aggregation functions was highlighted. The latter will mix texture properties of distinct tissues, yielding texture measurement that are not corresponding to anything visually or biologically (see Fig. 2.6). This was found to be widely overlooked in the literature and motivated the use of alternate aggregation functions (*e.g.*, covariances [10]) as well as defining digital tissue atlases of tumors or organs containing collections of regions for which it is reasonable to consider that texture properties are homogeneous. The latter provides the exciting opportunity to construct disease-specific digital phenotypes, which already showed to constitute powerful models for predicting disease diagnosis, treatment response and/or patient prognosis [12, 20, 19, 50] (see Fig. 2.8).

The second part of the chapter, Section 4, studied the type of directional information that is relevant for BTA. It was found that directional structures are defined locally but are not necessarily consistent over large regions. More precisely, a fundamental aspect of texture directionality is how directional structure intersect, which we called the Local Organization of Image Directions (LOID, see Section 4.1). The latter relate to the essential stitches of biomedical tissue, as well as to the texture primitives and texton theory discussed in Sections 2.4 of Chapter 1 and 2.3 of Chapter 3. Two adversarial categories of operators were analyzed in Section 4.2: directionally insensitive versus sensitive (see Fig. 2.12). Designing texture operators that are able to accurately characterize the LOIDs with robustness to rigid transformation was found to be challenging. On the one hand, directionally insensitive operators are invariant to rotations but are insensitive to image directions. On the other hand, directionally sensitive operators can sense directions but are not invariant to rotations. In addition, the effect of aggregation using integrative functions kills the ability of unidirectional operators to characterize the LOIDs even when the latter are all aligned to each other. In Section 4.3, we provided evidence that using Moving Frame (MF) texture representations consisting of locally aligning sets of non-collinear operators (see Eq. 2.5 and Fig. 2.14), allowed robust recognition of the LOIDs with invariance to rigid transformations. A quantitative comparison of the classification performance of texture classes with pronounced directional patterns and non-rigid transformations confirmed the superiority of MF representations, but at the expense of a high computational complexity when compared to much simpler directionally insensitive representations (see Section 4.4).

Overall, this chapter introduced a new set of comparison dimensions between BTA methods that is specific to biomedical imaging. The latter is further used in Chapter 3 to perform a systematic qualitative comparison of most popular BTA methods, which constitutes a user guide to assess the relevance of each approach for a particular

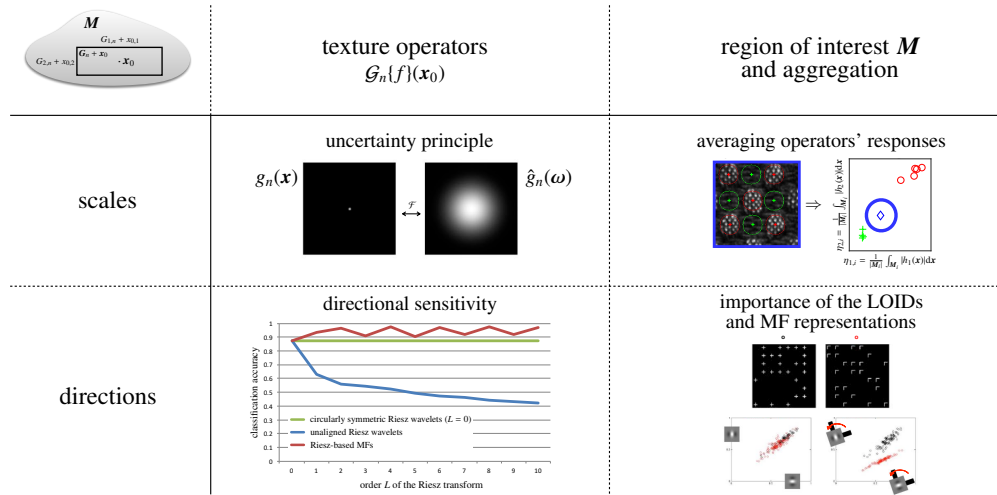


Figure 2.17: Checklist matrix of essential theoretic foundations and practical consequences in terms of choices of scales and directions for texture operator and aggregation function design.

medical or biological task in hand.

Acknowledgments

This work was supported by the Swiss National Science Foundation (under grant PZ00P2_154891), and the CIBM.

Bibliography

1. Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11):2274–2281, 2012.
2. Timo Ahonen, Ji? Matas, Chu He, and Matti Pietikinen. Rotation Invariant Image Description with Local Binary Pattern Histogram Fourier Features. In Arnt-Brrre Salberg, JonYngve Hardeberg, and Robert Jenssen, editors, *Image Analysis*, volume 5575 of *Lecture Notes in Computer Science*, pages 61–70. Springer Berlin Heidelberg, 2009.
3. Ohad Ben-Shahar and Steven W Zucker. The perceptual organization of texture flow: a contextual inference approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(4):401–417, apr 2003.
4. James R Bergen and Michael S Landy. Computational Modeling of Visual Texture Segregation. In *Computational Models of Visual Processing*, pages 253–271. MIT Press, 1991.
5. A.C. Bovik, M. Clark, and W.S. Geisler. Multichannel texture analysis using localized spatial filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(1):55–73, 1990.
6. Y Boykov, O Veksler, and R Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, nov 2001.
7. Alexandre Briguet, Isabelle Courdier-Fruh, Mark Foster, Thomas Meier, and Josef P Magyar. Histological parameters for the quantitative assessment of muscular dystrophy in the mdx-mouse. *Neuro-*

- muscular Disorders*, 14(10):675–682, 2004.
8. Emmanuel J Candès and David L Donoho. Curvelets – A Surprisingly Effective Nonadaptive Representation For Objects with Edges. In *Curves and Surface Fitting*, pages 105–120, Nashville, 2000. Vanderbilt University Press.
 9. Nicolas Chenouard and Michael Unser. 3D steerable wavelets and monogenic analysis for bioimaging. In *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 2132–2135, April 2011.
 10. Pol Cirujeda, Yashin Dicente Cid, Henning Müller, Daniel Rubin, Todd A Aguilera, Billy W Loo Jr., Maximilian Diehn, Xavier Binefa, and Adrien Depeursinge. A 3–D Riesz–Covariance Texture Model for Prediction of Nodule Recurrence in Lung CT. *IEEE Transactions on Medical Imaging*, 2016.
 11. Jia Deng, Wei Dong, R Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2009*, pages 248–255, 2009.
 12. Adrien Depeursinge, Anne C Chin, Ann N Leung, Donato Terrone, Michael Bristow, Glenn Rosen, and Daniel L Rubin. Automated Classification of Usual Interstitial Pneumonia using Regional Volumetric Texture Analysis in High-Resolution CT. *Investigative Radiology*, 50(4):261–267, apr 2015.
 13. Adrien Depeursinge, Antonio Foncubierta-Rodríguez, Dimitri Van De Ville, and Henning Müller. Rotation–covariant texture learning using steerable Riesz wavelets. *IEEE Transactions on Image Processing*, 23(2):898–908, February 2014.
 14. Adrien Depeursinge, Antonio Foncubierta-Rodríguez, Dimitri Van De Ville, and Henning Müller. Three-dimensional solid texture analysis and retrieval in biomedical imaging: review and opportunities. *Medical Image Analysis*, 18(1):176–196, 2014.
 15. Adrien Depeursinge, Pedram Pad, Anne C Chin, Ann N Leung, Daniel L Rubin, Henning Müller, and Michael Unser. Optimized steerable wavelets for texture analysis of lung tissue in 3-D CT: classification of usual interstitial pneumonia. In *IEEE 12th International Symposium on Biomedical Imaging, ISBI 2015*, pages 403–406. IEEE, apr 2015.
 16. Adrien Depeursinge, Zsuzsanna Püspöki, John-Paul Ward, and Michael Unser. Steerable wavelet machines (SWM): Learning moving frames for texture classification. *IEEE Transactions on Image Processing*, 26(4):1626–1636, 2017.
 17. Adrien Depeursinge, Dimitri Van De Ville, Alexandra Platon, Antoine Geissbuhler, Pierre-Alexandre Poletti, and Henning Müller. Near–affine–invariant texture learning for lung tissue analysis using isotropic wavelet frames. *IEEE Transactions on Information Technology in BioMedicine*, 16(4):665–675, July 2012.
 18. Adrien Depeursinge, Alejandro Vargas, Alexandra Platon, Antoine Geissbuhler, Pierre-Alexandre Poletti, and Henning Müller. Building a Reference Multimedia Database for Interstitial Lung Diseases. *Computerized Medical Imaging and Graphics*, 36(3):227–238, apr 2012.
 19. Adrien Depeursinge, Masahiro Yanagawa, Ann N Leung, and Daniel L Rubin. Predicting Adenocarcinoma Recurrence Using Computational Texture Models of Nodule Components in Lung CT. *Medical Physics*, 42:2054–2063, 2015.
 20. Yashin Dicente Cid, Henning Müller, Alexandra Platon, Jean-Paul Janssens, Lador Frédéric, Pierre-Alexandre Poletti, and Adrien Depeursinge. A Lung Graph-Model for Pulmonary Hypertension and Pulmonary Embolism Detection on DECT images. In *MICCAI Workshop on Medical Computer Vision: Algorithms for Big Data, MICCAI-MCV*, 2016.
 21. Yashin Dicente Cid, Henning Müller, Alexandra Platon, Pierre-Alexandre Poletti, and Adrien Depeursinge. Locally-Oriented Wavelet Transforms for 3-D Solid Texture Classification. *IEEE Transactions on Image Processing*, submitted.
 22. Olivier D Faugeras. Cartan’s Moving Frame Method and Its Application to the Geometry and Evolution of Curves in the Euclidean, Affine and Projective Planes. Technical report, Institut National de Recherche en Informatique et en Automatique (INRIA), 1993.
 23. Mary M Galloway. Texture analysis using gray level run lengths. *Computer Graphics and Image Processing*, 4(2):172–179, 1975.
 24. Mehrdad J Gangeh, Ali Ghodsi, and Mohamed S Kamel. Dictionary Learning in Texture Classifica-

- tion. In *Proceedings of the 8th international conference on Image analysis and recognition - Volume Part I*, pages 335–343, 2011.
25. Robert A Gatenby, Olya Grove, and Robert J Gillies. Quantitative Imaging in Cancer Evolution and Ecology. *Radiology*, 269(1):8–14, 2013.
 26. Paolo Gattuso, Vijaya B Reddy, Odile David, Daniel J Spitz, and Meryl H Haber. *Differential diagnosis in surgical pathology*. Elsevier Health Sciences, 2009.
 27. Diego Marcos Gonzalez, Michele Volpi, and Devis Tuia. Learning rotation invariant convolutional filters for texture classification. *CoRR*, abs/1604.0, 2016.
 28. Zhenhua Guo, Lei Zhang, and David Zhang. A Completed Modeling of Local Binary Pattern Operator for Texture Classification. *IEEE Transactions on Image Processing*, 19(6):1657–1663, jun 2010.
 29. Robert M Haralick. Statistical and Structural Approaches to Texture. *Proceedings of the IEEE*, 67(5):786–804, may 1979.
 30. Anil K Jain and Kalle Karu. Learning texture discrimination masks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(2):195–205, feb 1996.
 31. Svetlana Lazebnik, Cordelia Schmid, and Jean Ponce. A Sparse Texture Representation Using Local Affine Regions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1265–1278, aug 2005.
 32. Kun Liu, Henrik Skibbe, Thorsten Schmidt, Thomas Blein, Klaus Palme, Thomas Brox, and Olaf Ronneberger. Rotation-Invariant HOG Descriptors Using Fourier Analysis in Polar and Spherical Coordinates. *International Journal of Computer Vision*, 106(3):342–364, 2014.
 33. L Liu, S Lao, P W Fieguth, Y Guo, X Wang, and M Pietikinen. Median Robust Extended Local Binary Pattern for Texture Classification. *IEEE Transactions on Image Processing*, 25(3):1368–1381, mar 2016.
 34. D G Lowe. Object recognition from local scale invariant features. In *Proceedings of the International Conference of Computer Vision*, Corfu, Greece, sep 1999.
 35. David G Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
 36. Julien Mairal, Francis Bach, J Ponce, Guillermo Sapiro, and Andrew Zisserman. Supervised dictionary learning. *Advances in Neural Information Processing Systems*, pages 1033–1040, 2008.
 37. Timo Ojala, Topi Mäenpää, Matti Pietikäinen, Jaakko Viertola, Juha Kyllönen, and Sami Huovinen. Outex – new framework for empirical evaluation of texture analysis algorithms. In *16th International Conference on Pattern Recognition*, volume 1 of *ICPR*, pages 701–706. IEEE Computer Society, aug 2002.
 38. Timo Ojala, Matti Pietikäinen, and Topi Mäenpää. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(7):971–987, jul 2002.
 39. Maria Petrou. Texture in biomedical images. In Thomas M. Deserno, editor, *Biomedical Image Processing*, pages 157–176. Springer-Verlag Berlin Heidelberg, 2011.
 40. Maria Petrou and Pedro García Sevilla. *Non-Stationary Grey Texture Images*, pages 297–606. John Wiley & Sons, Ltd, 2006.
 41. Javier Portilla and Eero P Simoncelli. A Parametric Texture Model Based on Joint Statistics of Complex Wavelet Coefficients. *International Journal of Computer Vision*, 40(1):49–70, 2000.
 42. Z Püspöki, V Uhlmann, C Vonesch, and M Unser. Design of Steerable Wavelets to Detect Multifold Junctions. *IEEE Transactions on Image Processing*, 25(2):643–657, February 2016.
 43. Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. CNN features off-the-shelf: An astounding baseline for recognition. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 512–519, 2014.
 44. Laurent Sifre and Stéphane Mallat. Rigid-Motion Scattering for Texture Classification. *Submitted to International Journal of Computer Vision*, abs/1403.1:1–19, 2014.
 45. Eero P Simoncelli and William T Freeman. The steerable pyramid: a flexible architecture for multi-scale derivative computation. In *Proceedings of International Conference on Image Processing, 1995.*, volume 3, pages 444–447, oct 1995.

62 BIBLIOGRAPHY

46. L Sorensen, M Nielsen, P Lo, H Ashraf, J H Pedersen, and M de Bruijne. Texture-Based Analysis of COPD: A Data-Driven Approach. *IEEE Transactions on Medical Imaging*, 31(1):70–78, jan 2012.
47. M. Storath, A. Weinmann, and M. Unser. Unsupervised texture segmentation using monogenic curvelets and the Potts model. In *IEEE International Conference on Image Processing (ICIP)*, pages 4348–4352, October 2014.
48. Bart M ter Haar Romeny. Multi-Scale and Multi-Orientation Medical Image Analysis. In Thomas M Deserno, editor, *Biomedical Image Processing*, pages 177–196. Springer-Verlag Berlin Heidelberg, 2011.
49. Guillaume Thibault, Bernard Fertil, Claire Navarro, Sandrine Pereira, Pierre Cau, Nicolas Levy, Jean Sequeira, and Jean-luc Mari. Texture Indexes and Gray Level Size Zone Matrix Application to Cell Nuclei Classification. *Pattern Recognition and Information Processing*, pages 140–145, 2009.
50. Salvatore Torquato. Toward an Ising model of cancer and beyond. *Physical Biology*, 8(1):15017, 2011.
51. Michael Unser and Nicolas Chenouard. A unifying parametric framework for 2D steerable wavelet transforms. *SIAM Journal on Imaging Sciences*, 6(1):102–135, 2013.
52. Michael Unser, Nicolas Chenouard, and Dimitri Van De Ville. Steerable Pyramids and Tight Wavelet Frames in $L_2(\mathbb{R}^d)$. *IEEE Transactions on Image Processing*, 20(10):2705–2721, oct 2011.
53. Dimitri Van De Ville, Thierry Blu, and Michael Unser. Isotropic polyharmonic B-Splines: Scaling functions and wavelets. *IEEE Transactions on Image Processing*, 14(11):1798–1813, November 2005.
54. Manik Varma and Andrew Zisserman. A Statistical Approach to Texture Classification from Single Images. *International Journal of Computer Vision*, 62(1-2):61–81, 2005.
55. Hans Wang, Qing and Ronneberger, Olaf and Burkhardt. Fourier analysis in polar and spherical coordinates. Technical report, Albert-Ludwigs-Universität Freiburg, Institut für Informatik, 2008.
56. J P Ward and M Unser. Harmonic Singular Integrals and Steerable Wavelets in $L_2(\mathbb{R}^d)$. *Applied and Computational Harmonic Analysis*, 36(2):183–197, mar 2014.
57. Dong-Hui Xu, Arati S Kurani, Jacob Furst, and Daniela S Raicu. Run-length encoding for volumetric texture. In *The 4th IASTED International Conference on Visualization, Imaging, and Image Processing – VIIP 2004*, Marbella, Spain, sep 2004.
58. Y Xu, X Yang, H Ling, and H Ji. A new texture descriptor using multifractal analysis in multi-orientation wavelet pyramid. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 161–168, jun 2010.