

Bayesian Multifractal Image Segmentation

Kareth M. León-López¹, Abderrahim Halimi²,
Jean-Yves Tourneret¹, and Herwig Wendt^{1*}

¹ IRIT Laboratory, Université de Toulouse, CNRS, INP-Toulouse, UT3, UT2, TésA, Toulouse, France.

{firstname.lastname}@irit.fr

² Heriot-Watt University, Edinburgh EH14 4AS, UK.

a.halimi@hw.ac.uk

Abstract

Multifractal analysis (MFA) provides a framework for the global characterization of image textures by describing the spatial fluctuations of their local regularity based on the multifractal spectrum. Several works have shown the interest of using MFA for the description of homogeneous textures in images. Nevertheless, natural images can be composed of several textures and, in turn, multifractal properties associated with those textures. This paper introduces an unsupervised Bayesian multifractal segmentation method to model and segment multifractal textures by jointly estimating the multifractal parameters and labels on images, at the pixel-level. For this, a computationally and statistically efficient multifractal parameter estimation model for wavelet leaders is firstly developed, defining different multifractality parameters for different regions of an image. Then, a multiscale Potts Markov random field is introduced as a prior to model the inherent spatial and scale correlations (referred to as cross-scale correlations) between the labels of the wavelet leaders. A Gibbs sampling methodology is finally used to draw samples from the posterior distribution of the unknown model parameters. Numerical experiments are conducted on synthetic multifractal images to evaluate the performance of the proposed segmentation approach. The proposed method achieves superior performance compared to traditional unsupervised segmentation techniques as well as modern deep learning-based approaches, showing its effectiveness for multifractal image segmentation.

Keywords— Texture segmentation, wavelet leaders, multifractal analysis, Bayesian estimation, Gibbs sampling, multiscale granularity coefficient, Potts Markov random field.

1 Introduction

Scaling, or scale invariance, is a modeling paradigm in which a given object structure or statistical property remains unchanged regardless of the scale at which it is analyzed [1], and can be exploited using self-similar or multifractal models [2]. The scaling properties of texture images have been studied via the multifractal analysis (MFA) that allows complex patterns (local transient regularity fluctuations) to be described and modeled across different scales [3, 4]. MFA has been exploited in several applications of different natures, including network traffic analysis [5], texture classification [6, 7, 8], texture characterization [9] and art analysis [1]. Current state-of-the-art MFA tools rely on the use of dedicated multiresolution coefficients, the wavelet leaders [2, 7, 10].

On the other hand, texture segmentation is a long-standing problem in image processing that has received significant research efforts in the past and continues to attract considerable attention in the scientific community [11, 12, 13]. Several methods have relied on geometrical or statistical features using Markov random fields [14, 15], while others rely on pixel intensities to group different regions [16, 17]. Other works have exploited the scale invariance property for image segmentation [18, 19]. As an example, self-similarity exponents and local regularity have been used to segment image textures [18]. Works such

*This work was partly supported by Grant ANR-18-CE45-0007 MUTATION and the UK Royal Academy of Engineering under the Research Fellowship Scheme (RF/201718/17128).

as [20] have also used fractal information to extract features. The literature reports some MFA-based image segmentation methods that rely on the multifractal spectrum [19, 21, 22] (see definition in Section II-A and see a review of the methods in [23]). For instance, the multifractal spectrum was used in [21] to segment water bodies in synthetic aperture radar (SAR) images. However, the strategy investigated in [21] is computationally demanding and relies on an empirical threshold for region delineation. Segmentation approaches based on multifractal descriptors have also been studied [24, 25]. The state-of-the-art MFA formalism (in terms of precision and computational efficiency) is based on wavelet leaders [26, 4, 27]. In this line, authors in [4] provide a methodology for modeling the multifractality of multivariate images, yielding an operational tool for the multifractal analysis of multidimensional images as well as images with several multifractal textures. However, the method is not intended to assess multifractality at the pixel level or for segmentation purposes, which is a key limitation of current multifractal-based approaches. Indeed, characterizing point-wise regularity fluctuations requires to analyze several scales, tied to space averages [28], thus yielding a paradox and an important difficulty for pixel-level analysis and segmentation.

Multi-resolution coefficients associated with an image have been used in several works to perform image segmentation. The wavelet transform was used in [29] and [30] to compute multi-resolution coefficients of an image that are combined from coarse to fine scales for image segmentation. Further works such as [31, 32] have proposed multiscale Bayesian techniques in which a contextual behavior from a coarse scale (low-resolution scale) is employed up to the finest scale (highest resolution) to provide an image segmentation. Multiscale segmentation has also been performed using hidden Markov trees [29, 33] or neural networks [34]. The main limitation of these approaches is that they rely on training data to describe the texture parameters of the images. Additionally, even though there exist works in the literature that attempt to consider coefficients at multiple scales for image segmentation, up to our knowledge, there is no method in the state of the art that exploits these scales to segment images at the pixel level based on a multifractal model.

This work introduces a new multifractal image segmentation approach for the joint parameter and label estimation of heterogeneous images via a multiscale decomposition and a Bayesian framework. The proposed segmentation approach is fully unsupervised and is conducted in the wavelet leader domain, in which the scale-free textures are parameterized by the multifractality parameter. The contributions of this paper are the following:

- First, a computationally efficient Fourier-based Whittle approximation for the multifractality parameter estimation on heterogeneous multifractal images is introduced. This contribution relies on the efficient model developed in [4] that is extended and adapted to *irregular shapes* for modeling the different piecewise textures of an image.
- Second, a likelihood function defined in terms of the *multifractality parameter* across scales is introduced, leading to a model for scale invariant textures.
- Third, a multiscale Potts Markov random field that considers the *intra-scale* and *inter-scale* dependencies in the wavelet leader domain is considered. This random field uses a multiscale graph structure that exploits spatial and scale image correlations. Moreover, the granularity coefficients of the proposed multiscale Potts-Markov random field are efficiently and jointly sampled with the labels.
- Fourth, we conduct extensive numerical experiments and comparisons with state-of-the-art image segmentation methods for simulated and real-world images, showing the interest of the proposed framework.

2 Multifractal Analysis of Images

2.1 Local Regularity and Multifractal Spectrum

MFA enables the study of local regularity of functions. To that end, it makes use of a quantity that measures the so-called Hölder exponent $h(\mathbf{t})$ at a location \mathbf{t} , and a *global* description of its spatial fluctuations, the multifractal spectrum $\mathcal{D}(h)$ [35].

Hölder Exponent. Let $\mathbf{F}(\mathbf{t}) : \mathbb{R}^2 \rightarrow \mathbb{R}$ denote the function (image) to be analyzed, under the assumption that \mathbf{F} is locally bounded. Then, \mathbf{F} is said to belong to $C^\alpha(\mathbf{t}_0)$ with $\alpha > 0$ if there exists a positive constant C and a polynomial $P_{\mathbf{t}_0}$, of degree smaller than α , such that $\|\mathbf{F}(\mathbf{t}) - P_{\mathbf{t}_0}(\mathbf{t})\| \leq C\|\mathbf{t} - \mathbf{t}_0\|^\alpha$ holds, where $\|\cdot\|$ is the Euclidean norm. The Hölder exponent at \mathbf{t}_0 is then defined as

$$h(\mathbf{t}_0) \triangleq \sup\{\alpha : \mathbf{F} \in C^\alpha(\mathbf{t}_0)\}.$$

Here, large (small) values of $h(\mathbf{t}_0)$ correspond to locations where \mathbf{F} is locally smooth (irregular), respectively [18].

Multifractal Spectrum. The collection of fractal dimensions (Hausdorff dimension \dim_H) of the sets of points for which the Hölder exponent takes the same value defines the multifractal spectrum

$$\mathcal{D}(h) \triangleq \dim_H \{\mathbf{t} : h(\mathbf{t}) = h\}. \quad (1)$$

The estimation of $\mathcal{D}(h)$ is the main target of MFA. However, it is unfeasible using its formal definition (1). Indeed, the estimation relies on a multiscale decomposition and the so-called *multifractal formalism* (for more details, see [2, 4, 10]). The current state-of-the-art multifractal formalism is based on the wavelet leaders of the two-dimensional (2D) discrete wavelet transform (DWT) coefficients of the given image and is recalled next.

2.2 Multifractal Formalism based on Wavelet Leaders

Wavelet Leaders. The 2D orthonormal DWT of $\mathbf{F} \in \mathbb{R}^{N \times N}$ is obtained by considering four 2D filters $\phi^{(m)}(\mathbf{n})$, $m = 0, \dots, 3$, defined as the tensor products of a low-pass filter $\phi(n)$ (scaling function) and a high-pass filter $\psi(n)$ (wavelet function characterized by $N_\psi > 0$ vanishing moments). Precisely, the four 2D filters are written as [27, 4]:

$$\begin{cases} \psi^{(0)}(\mathbf{n}) = \phi(n_1)\phi(n_2), & \psi^{(1)}(\mathbf{n}) = \psi(n_1)\phi(n_2), \\ \psi^{(2)}(\mathbf{n}) = \phi(n_1)\psi(n_2), & \psi^{(3)}(\mathbf{n}) = \psi(n_1)\psi(n_2), \end{cases}$$

where $\psi^{(0)}(\mathbf{n})$ is the 2D low-pass filter and $\{\psi^{(m)}(\mathbf{n})\}_{m=1}^3$ are the 2D high-pass filters. Let $D_{\mathbf{F}}^{(m)}(j, \mathbf{n})$ denotes the 2D DWT of \mathbf{F} , with $D_{\mathbf{F}}^{(0)}(0, \mathbf{n}) \triangleq \mathbf{F}$. The DWT coefficients $\{D_{\mathbf{F}}^{(m)}(j, \mathbf{n})\}_{m=1}^3$, and the approximation coefficients $D_{\mathbf{F}}^{(0)}(j, \mathbf{n})$, for $j \geq 1$, are obtained by iteratively convolving $\psi^{(m)}$ with the coefficients $D_{\mathbf{F}}^{(0)}(j-1, \cdot)$, followed by decimation [4]. For MFA purposes, the 2D DWT coefficients are normalized as $d_{\mathbf{F}}^{(m)}(j, \mathbf{n}) \triangleq 2^{-j} D_{\mathbf{F}}^{(m)}(j, \mathbf{n})$, for $m = 1, 2, 3$.

The *wavelet leaders* are defined in [26] as the largest DWT coefficient magnitudes within the spatial neighborhood $3\lambda_{j,\mathbf{n}}$ over all finer scales:

$$L_{j,\mathbf{n}} = \sup_{\substack{m \in \{1, 2, 3\} \\ \lambda' \subset 3\lambda_{j,\mathbf{n}}}} |d_{\lambda'}^{(m)}|, \quad (2)$$

where $\lambda_{j,\mathbf{n}}$ is the dyadic cube of side length 2^j at position $2^j \mathbf{n}$, i.e., $\lambda_{j,\mathbf{n}} = \{[2^j n_1, 2^j(n_1+1)), [2^j n_2, 2^j(n_2+1))\}$ [36, 37], and $3\lambda_{j,\mathbf{n}} = \bigcup_{i_1, i_2 \in \{-1, 0, 1\}} \lambda_{j, (n_1+i_1, n_2+i_2)}$ (see [26] for more details about wavelet leaders and related bounds).

It has been shown that the coefficients c_p are tied to the wavelet leaders $L_{j,\mathbf{n}}$ via the key relation [27]

$$\text{Cum}_p[\ln L_{j,\mathbf{n}}] = c_p^0 + c_p \ln 2^j. \quad (3)$$

In addition, the multifractal formalism provides an expansion of $\mathcal{D}(h)$ in terms of the log-cumulants c_p , $p \geq 1$, such that

$$\mathcal{D}(h) = 2 + \frac{c_2}{2!} \left(\frac{h - c_1}{c_2} \right)^2 + \frac{-c_3}{3!} \left(\frac{h - c_2}{c_2} \right)^3 + \dots, \quad (4)$$

for $c_2 < 0$, see [4, 27] for further details. These so-called log-cumulants provide a relevant summary of $\mathcal{D}(h)$, e.g., c_1 is the mode of $\mathcal{D}(h)$ that can be interpreted as the average smoothness of the image \mathbf{F} and is related to the Hurst parameter [38], c_2 is related to the width of $\mathcal{D}(h)$ and captures the degree of fluctuations of the local regularity of \mathbf{F} and is the most important feature of multifractal functions, termed *multifractality parameter*. Figure 1 presents a schematic illustration of the theoretical multifractal spectrum for two homogeneous 2D multifractal random walks, shown for $c_1 = 0.5, c_2 = -0.02$ (blue curve) and $c_1 = 0.98, c_2 = -0.08$ (black curve) [2, 3, 7] (note that these are the multifractal spectra of the composed image of Fig. 4-a-top).

Estimation of c_2 . The relation (3) with $p = 2$ leads to $\text{Var}[\ln L_{j,\mathbf{n}}] = c_2^0 + c_2 \ln 2^j$, which indicates that c_2 can be estimated by linear regression over a range of scale $j \in [j_1, j_2]$:

$$\hat{c}_2 = \frac{1}{\ln 2} \sum_{j=j_1}^{j_2} w_j \widehat{\text{Var}}[\ln(L_{j,\mathbf{n}})], \quad (5)$$

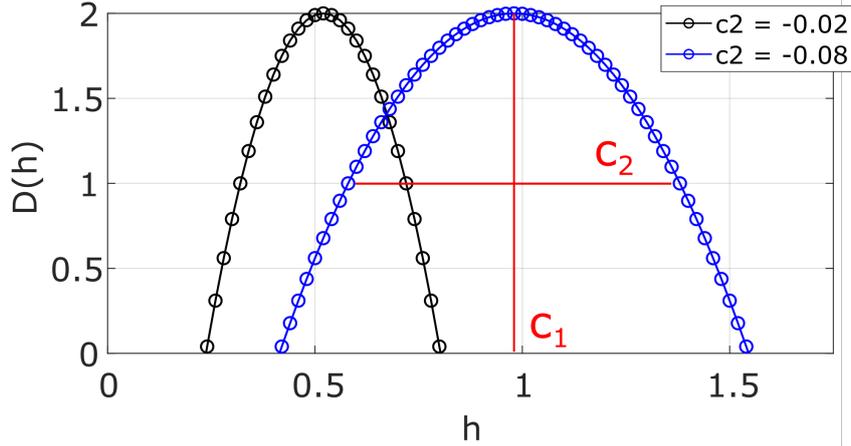


Figure 1: Multifractal spectrum $\mathcal{D}(h)$ schematic representation from two 2D multifractal random walk images, where c_1 is the mode and c_2 , the width, is the multifractality parameter; (c_1, c_2) capture most of the local regularity information of the image.

where $\widehat{\text{Var}}(\cdot)$ denotes the sample variance and w_j are suitable regression weights [2, 3, 39]. The estimator (5) has low complexity but poor performance (large variance). To overcome this limitation, a Fourier-based Bayesian model was introduced for homogeneous images [4]. It is well suited for the design of joint priors and briefly recalled next.

2.3 Estimation of c_2 using a Fourier-based Bayesian model

Let $\ell_{j,\mathbf{n}} = \ln(L_{j,\mathbf{n}})$ denote the centered logarithms of the wavelet leader coefficients, referred to as *log-leaders*, at scale j and at position \mathbf{n} . The vector of log-leaders, $\ell_j \in \mathbb{R}^{N_j^2}$, is obtained by vectorization using the lexicographic order, where $N_j = N/2^j$. Note that $j \in \llbracket j_1, j_2 \rrbracket$, $1 \leq j_1 \leq \dots \leq j_2 \leq J$, in which $j = 1$ and $j = J$ index the finest and the coarsest scales, respectively, where $\llbracket a, b \rrbracket$ denotes the set of integers ranging from a to b . Further, let $\ell \triangleq [\ell_{j_1}^T, \dots, \ell_{j_2}^T]^T$ be the vector obtained by concatenation of the log-leaders at scales j_1, \dots, j_2 . According to [4, 27], the distribution of the log-leaders ℓ_j can be well-approximated by a multivariate Gaussian distribution with density given by:

$$p(\ell_j | \theta_1, \theta_2) \propto |\Sigma_j(\theta)|^{-1/2} \exp\left(-\frac{1}{2} \ell_j^T \Sigma_j(\theta, \|\Delta\mathbf{n}\|)^{-1} \ell_j\right). \quad (6)$$

Assuming independence between scales, one obtains:

$$p(\ell | \theta_1, \theta_2) = \prod_{j=j_1}^{j_2} p(\ell_j | \theta_1, \theta_2), \quad (7)$$

where $|\mathbf{A}|$ denotes the determinant of a matrix \mathbf{A} , $\theta \triangleq [\theta_1, \theta_2]^T$, $-c_2 = \theta_1$, θ_2 is related to the model adjustment constant c_2^0 , and $\Sigma_j(\theta, \|\Delta\mathbf{n}\|) \approx \text{Cov}(\ell_{j,\mathbf{n}}, \ell_{j,\mathbf{n}+\Delta\mathbf{n}})$ is a covariance model [27], defined by the radial symmetric function

$$\Sigma_j(\theta, r) \triangleq \theta_1 g_1(j, r) + \theta_2 g_2(j, r), \quad (8)$$

where $r \triangleq \|\Delta\mathbf{n}\|$, g_1 and g_2 are functions defined as $g_1(j, r) = \max\{0, -\ln((r+1)/(r_j+1))\}$ and $g_2(j, r) = \max\{0, -\ln(r+1)/\ln 4\}$ with $r_j = \lfloor N_j/4 \rfloor$ [40].

The likelihood (7) is difficult to be used in a numerical algorithm but can be efficiently approximated in the spectral domain by formulating an equivalent Fourier-domain data augmented model introduced in [4]. This statistical model is based on a Whittle approximation and data augmentation allowing the suitable design of conjugate priors for the parameters θ_1, θ_2 [4, 27]. More precisely, (7) is approximated as

$$p(\ell | \theta_1, \theta_2) \approx p(\mathbf{x}, \boldsymbol{\mu} | \theta_1, \theta_2), \quad (9)$$

where $p(\mathbf{x}, \boldsymbol{\mu} | \theta_1, \theta_2)$ is given by

$$p(\mathbf{x}, \boldsymbol{\mu} | \theta_1, \theta_2) \propto \theta_1^{-S} \exp\left(-\frac{1}{\theta_1} (\mathbf{x} - \boldsymbol{\mu})^H G_{1,s}^{-1} (\mathbf{x} - \boldsymbol{\mu})\right) \theta_2^{-S} \exp\left(-\frac{1}{\theta_2} \boldsymbol{\mu}^H G_{2,s}^{-1} \boldsymbol{\mu}\right), \quad (10)$$

where $\boldsymbol{\mu}$ is a latent vector and $\mathbf{x} = [\mathbf{x}_1^T, \dots, \mathbf{x}_S^T]^T$ is a subset of Fourier coefficients of ℓ [4, 40],

$$\mathbf{x}_s \triangleq \mathbf{x}(\boldsymbol{\omega}_s) = \frac{1}{N_j} \sum_{\mathbf{n} \in \llbracket -N_j, N_j \rrbracket^2} \ell_{j,\mathbf{n}} \exp(-i \mathbf{n}^T \boldsymbol{\omega}_s), \quad (11)$$

where $s \in \{(j, \mathbf{m}) : j = \llbracket j_1, j_2 \rrbracket, \mathbf{m} \in \Omega_j^\dagger\}$, Ω_j^\dagger indexes low frequencies (see [4, 40] for further details), and $\boldsymbol{\omega}_s = 2\pi\mathbf{m}/\sqrt{N_j}$ is the given Fourier frequency, with $(j, \mathbf{m}) \in s$. Moreover, the matrices

$$G_{i,s} \triangleq \sum_{\mathbf{n} \in \llbracket -N_j, N_j \rrbracket^2} g_i(j, \|\mathbf{n}\|) \exp(-i \mathbf{n}^T \boldsymbol{\omega}_s), \quad (12)$$

for $i = 1, 2$ are real and strictly positive matrices defined on s . The corresponding spectral covariance of the asymptotic approximation (10) is expressed as

$$\boldsymbol{\Xi}_s = \theta_1 G_{1,s} + \theta_2 G_{2,s}, \quad (13)$$

which leads to standard conditional distributions and fast estimation when used in a Bayesian model [4, 40].

3 Towards Multifractal-based Image Segmentation in the Fourier Domain

Model (10) provides an accurate estimation of c_2 (recall that $-c_2 = \theta_1$). However, it is not designed for images with several multifractal textures. Here, we propose a new statistical model in the Fourier domain for modeling heterogeneous multifractal images. The limitation of (10) for computing c_2 is its reliance on a regular spatial grid. This prohibits its use for scenarios involving heterogeneous image regions with irregular shapes having different multifractal properties. We therefore develop a spectral likelihood that can handle multifractal regions with irregular shapes [41, 42], resulting in a more flexible model. It is important to note that this spectral likelihood is developed by means of the debiased spatial Whittle likelihood from [42], used and adapted to address the complexity of computing the likelihood in the Fourier domain for irregularly shaped regions on regularly sampled images. Although the approach could be adapted to handle truly irregularly sampled data, this is beyond the scope of this work.

3.1 Likelihood on Regions With Irregular Shapes

Consider an image \mathbf{F} that is composed of regions with different multifractal spectra. The log-leaders of \mathbf{F} can be assumed to be a mixture of K classes denoted as \mathcal{C}_k with different multifractal parameters $\boldsymbol{\theta}_k \triangleq [\theta_1^k, \theta_2^k]^T$, $k = 1, \dots, K$. For segmentation purposes, let $z_{j,n}$ be the label at position n and scale j that maps the observation $\ell_{j,n}$ to the class \mathcal{C}_k . The label vector at scale j is denoted as $\mathbf{z}_j = [z_{j,1}, \dots, z_{j,N_j^2}]^T$ and $\mathbf{z} = [\mathbf{z}_{j_1}^T, \dots, \mathbf{z}_{j_2}^T]^T$. Let $I_{j,k} = \{n \in N_j^2 | z_{j,n} = k\}$ be the subset of sites in \mathbf{z}_j with label k at scale j , and $\mathcal{T}_{j,\mathbf{n},k}$ its corresponding indicator function:

$$\mathcal{T}_{j,\mathbf{n},k} = \begin{cases} 1, & \mathbf{n} \in I_{j,k} \\ 0, & \text{otherwise.} \end{cases} \quad (14)$$

The likelihood (7) can then be rewritten as

$$p(\ell | \boldsymbol{\Theta}, \mathbf{z}) = \prod_{j=j_1}^{j_2} \prod_{k=1}^K p(\ell_{j,\mathbf{n} \in I_{j,k}} | \boldsymbol{\theta}_k^T), \quad (15)$$

where $p(\ell_{j,\mathbf{n} \in I_{j,k}} | \boldsymbol{\theta}_k^T)$ is the likelihood of a stationary zero mean Gaussian process with covariance (8) observed at locations $\mathbf{n} \in I_{j,k}$, and $\boldsymbol{\Theta} = [\boldsymbol{\theta}_1^T, \dots, \boldsymbol{\theta}_K^T]^T$ contains the whole set of multifractal parameters. Just like (7), (15) is problematic for direct use in practice. In addition, obtaining an approximation similar to (10) for (15) is less straightforward.

3.1.1 Debiased Fourier Coefficients

Having different multifractality regions implies rewriting (10) using the coefficients that share the same multifractality, yielding an irregular shape which requires to adjust for bias in the Fourier transform [41, 42]. The log-leader Fourier coefficients are redefined for the k -th class label in terms of the irregular domain $\mathcal{T}_{j,n,k}$ as

$$\mathbf{x}_{s,k} \triangleq \xi_{j,k} \sum_{\mathbf{n}} \mathcal{T}_{j,\mathbf{n},k} [\ell_{j,\mathbf{n}} - t_{j,k}] \exp(-i \mathbf{n}^T \boldsymbol{\omega}_s), \quad (16)$$

where $\xi_{j,k} = 1/(\sum_{\mathbf{n}} \mathcal{T}_{j,\mathbf{n},k}^{1/2})$ and $t_{j,k} = \xi_{j,k}^{-1} \sum_{\mathbf{n}} \mathcal{T}_{j,\mathbf{n},k} \ell_{j,\mathbf{n}}$ are scaling constants considering the bias introduced by the irregular domain [41, 42].

3.1.2 Covariance Matrix

The covariance matrix in (13) for the k -th class is rewritten as

$$\boldsymbol{\Xi}_{s,k} \triangleq \mathbf{W}_{s,k} \boldsymbol{\Xi}_s, \quad (17)$$

$$= \theta_1^k \mathbf{W}_{j,\mathbf{n},k} G_{1,s} + \theta_2^k \mathbf{W}_{j,\mathbf{n},k} G_{2,s}, \quad (18)$$

where $\mathbf{W}_{j,\mathbf{n},k}$ is a weight matrix defined as [41, 42]:

$$\mathbf{W}_{j,(n_1,n_2),k} = \xi_{j,k} \sum_{u \in \mathcal{R}_j} \sum_{r \in \mathcal{R}_j} \mathcal{T}_{j,(u,r),k} \mathcal{T}_{j,(u+n_1,r+n_2),k}, \quad (19)$$

where $\mathcal{R}_j = \llbracket -N_j, N_j - 1 \rrbracket^2$ is a regular grid. The likelihood of the k -th set is given by

$$p(\mathbf{x}_k | \theta_1^k, \theta_2^k, \mathbf{z}) \propto \prod_{s=1}^S |\boldsymbol{\Xi}_{s,k}|^{-1/2} \exp(-\mathbf{x}_{s,k}^H \boldsymbol{\Xi}_{s,k}^{-1} \mathbf{x}_{s,k}), \quad (20)$$

where $\mathbf{x}_k = [\mathbf{x}_{1,k}, \dots, \mathbf{x}_{S,k}]^T$ is the vector of Fourier coefficients of class k .

Accordingly, making use of data augmentation as in (10), the likelihood (20) can be written in a form that factorizes terms with the parameters θ_1^k and θ_2^k , as

$$p(\mathbf{x}_k, \boldsymbol{\mu}_k | \theta_1^k, \theta_2^k, \mathbf{z}) \propto \prod_{s=1}^S (\theta_2^k)^{-S} \exp\left(-\frac{1}{\theta_2^k} \boldsymbol{\mu}_k^H \mathbf{W}_{s,k} G_{2,s}^{-1} \boldsymbol{\mu}_k\right) (\theta_1^k)^{-S} \exp\left(-\frac{1}{\theta_1^k} (\mathbf{x}_k - \boldsymbol{\mu}_k)^H \mathbf{W}_{s,k} G_{1,s}^{-1} (\mathbf{x}_k - \boldsymbol{\mu}_k)\right), \quad (21)$$

where $\mathbf{W}_{s,k} \triangleq \mathbf{W}_{j,\mathbf{n},k}$, and $\boldsymbol{\mu}_k = [\mu_{1,k}^T, \dots, \mu_{S,k}^T]^T$ is the complex-valued latent vector. The likelihood for all classes is finally defined as:

$$p(\mathbf{x}, \boldsymbol{\mu} | \boldsymbol{\Theta}, \mathbf{z}) = \prod_{k=1}^K p(\mathbf{x}_k, \boldsymbol{\mu}_k | \boldsymbol{\theta}_k, \mathbf{z}), \quad (22)$$

where $\boldsymbol{\theta}_k = [\theta_1^k, \theta_2^k]^T$. Finally, as in (9), we use (22) as an approximation for (15):

$$p(\ell | \boldsymbol{\Theta}, \mathbf{z}) \approx p(\mathbf{x}, \boldsymbol{\mu} | \boldsymbol{\Theta}, \mathbf{z}). \quad (23)$$

4 Bayesian Model for Multifractal Segmentation

With the likelihood (22), the multifractal parameters and labels corresponding to image regions with different regularities can be estimated using Bayesian estimators. This approach uses Bayesian inference, where prior distributions are assigned to the unknown model parameters, and these parameters are estimated using the resulting posterior distribution. In this section, the proposed Bayesian model for image segmentation and multifractal parameter estimation is presented.

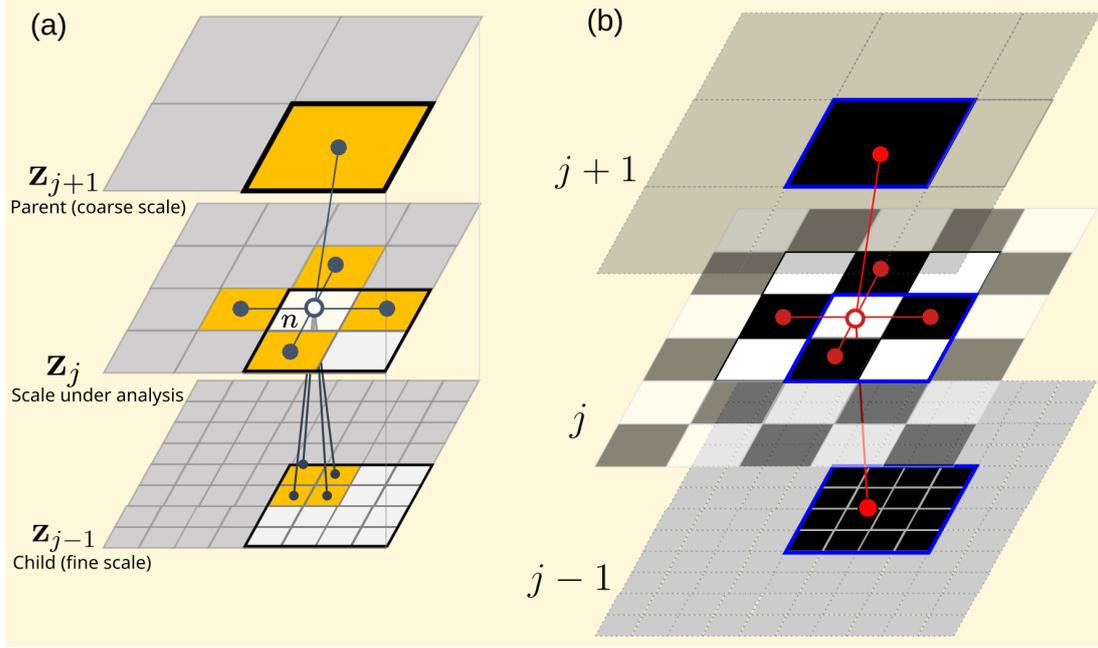


Figure 2: (a) Multiscale label graph structure considered as a prior on the log-leader coefficients. The label at position n (white circle) uses its spatial and scale neighbors (gray dots) to update its value accordingly. (b) Chessboard label decomposition for parallelizing the label sampling in the Potts-Markov random field, where the image at scale j is divided into two independent sets of pixels — white and black, and they are assumed independent. For each scale, pixels are updated using the chessboard scheme considering the spatio-scale connections marked in red.

4.0.1 Multiscale Label Prior

To account for the multiscale nature of the log-leaders, the label dependencies within and across scales are modeled via a graph. Specifically, *inter-scale* connections link neighbors across scales, and the *intra-scale* connections model the spatial correlation of labels. Figure 2-(a) illustrates the proposed graph structure at a given scale j , where the labels to be sampled at position n (white circle) use the values of its neighbors in space and scale (gray dots) to update their values accordingly. Formally, the corresponding multiscale prior for $\mathbf{z}_j \in \{1, \dots, K\}^n$ at scale j is defined as a modified version of the classical Potts-Markov field [43],

$$p(\mathbf{z}_j | \mathbf{z}_{-j}, \beta_s, \beta_{xy}^j) = \frac{1}{Z} \exp \left[\Phi_{\beta_s}(\mathbf{z}_j, \mathbf{z}_{-j}) + \Phi_{\beta_{xy}^j}(\mathbf{z}_j) \right], \quad (24)$$

where \mathbf{z}_{-j} denotes the labels \mathbf{z}_{j-1} , \mathbf{z}_{j+1} on adjacent scales (if $j = 1$ or $j = J$, then $\mathbf{z}_{-j} = \mathbf{z}_2$ or $\mathbf{z}_{-j} = \mathbf{z}_{J-1}$, respectively).

The potential for the *spatial* (intra-scale) dependencies is:

$$\Phi_{\beta_{xy}^j}(\mathbf{z}_j) \triangleq \sum_{n=1}^{N_j^2} \sum_{m \in \mathcal{V}(n)} \beta_{xy}^j \delta(z_{j,n} - z_{j,m}), \quad (25)$$

where $\mathcal{V}(n)$ contains the four spatial neighbors of n [44], as depicted in Fig. 2-(a). The potential for the *scale* (inter-scale) dependencies is defined as:

$$\Phi_{\beta_s}(\mathbf{z}_j, \mathbf{z}_{-j}) \triangleq \sum_{n=1}^{N_j^2} \sum_{p \in D(n)} \beta_s \delta(z_{j,n} - z_{j+1,p}) + \sum_{q \in U(n)} \beta_s \delta(z_{j,n} - z_{j-1,q}), \quad (26)$$

where $D(n)$ and $U(n)$ are multiscale index mapping functions that map the index n to its corresponding parent (*down*) or child (*up*) coefficient at a given scale, respectively (see Fig. 2-(a)). The inter-scale and spatial granularity coefficients β_s and β_{xy}^j , respectively, adjust the degree of homogeneity across scales

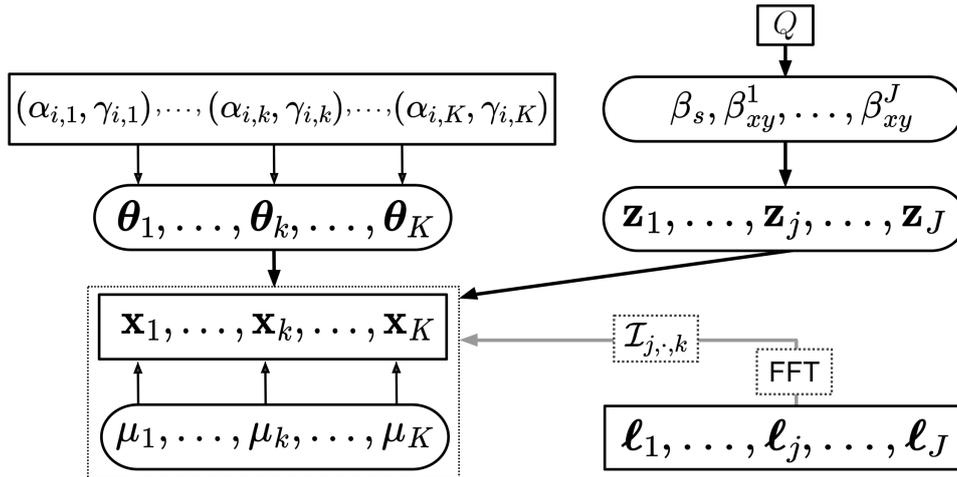


Figure 3: Directed acyclic graph (DAG) of the Bayesian model: the ellipse nodes show the estimated quantities and the square ones, the fixed ones. The FFT transform links the ℓ parameters to \mathbf{x} and $\boldsymbol{\mu}$.

and in space, and are gathered in the vector $\boldsymbol{\beta} = [\beta_s, \beta_{xy}^1, \dots, \beta_{xy}^J]$. Finally, $\tilde{C} = C(\beta_s, \beta_{xy}^j)$ is a partition function defined as

$$C(\beta_s, \beta_{xy}^j) \triangleq \sum_{\mathbf{z}_j} \exp [\Phi_{\beta_s}(\mathbf{z}^j, \mathbf{z}_{-j}) + \Phi_{\beta_{xy}^j}(\mathbf{z}_j)], \quad (27)$$

which is generally intractable because the number of summation terms grows exponentially with the size of \mathbf{z}_j [45].

4.1 Parameter Priors

4.1.1 Granularity Parameters β_s and β_{xy}^j

Fixing a specific value for the parameters β_s and β_{xy}^j in (24) is a difficult task due to individual image's spatial structure and organization that can change from one image to another. Indeed, small values of β_{xy}^j lead to noisy segmentation whereas large ones lead to coarse or over-smoothed regions. Inspired by [46, 45], the granularity parameters are jointly estimated with the labels and are assigned a uniform prior distribution on the interval $(0, Q)$, i.e., $f(\boldsymbol{\beta}) = \mathcal{U}_{(0,Q)}(\boldsymbol{\beta})$, where $\boldsymbol{\beta} = [\beta_s, \beta_{xy}^1, \dots, \beta_{xy}^J]$ and Q is a given constant related to the number of classes K .

4.1.2 Multifractal Parameter Prior

The natural conjugate priors for the parameters $\boldsymbol{\theta}_i^k$ are inverse-gamma distributions $\mathcal{IG}(\alpha_{i,k}, \gamma_{i,k})$ [4, 27], with fixed constants $\alpha_{i,k}, \gamma_{i,k}$ ($i = 1, 2$) defined, here, for each multifractal region $k = 1, \dots, K$.

4.2 Posterior Distribution

The posterior distribution of the model parameters depends on the spectral likelihood (22) and the priors as follows:

$$p(\boldsymbol{\Theta}, \boldsymbol{\mu}, \mathbf{z}, \boldsymbol{\beta} | \mathbf{x}) \propto p(\mathbf{x}, \boldsymbol{\mu} | \boldsymbol{\Theta}, \mathbf{z}) p(\boldsymbol{\Theta}) p(\mathbf{z} | \boldsymbol{\beta}) p(\boldsymbol{\beta}). \quad (28)$$

The Bayesian model is summarized in Fig. 3. Note that given the complexity of the posterior distribution (28), closed-form expressions for the estimators of the unknown parameters $\boldsymbol{\Theta}$, \mathbf{z} and $\boldsymbol{\beta}$ cannot be obtained. As an alternative, it is possible to generate samples that are asymptotically distributed according to the target distribution (28) [45]. These samples are then used to approximate Bayesian estimators of the unknown model parameters.

5 Sampling Procedure

This section describes a Gibbs sampler that draws samples distributed according to the posterior distribution (28) by making use of its conditional distributions.

5.1 Conditional Probability $p(\mathbf{z}|\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Theta}, \boldsymbol{\beta})$

The conditional distribution for all labels \mathbf{z} is

$$p(\mathbf{z}|\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Theta}, \boldsymbol{\beta}) \propto p(\mathbf{x}, \boldsymbol{\mu}|\boldsymbol{\Theta}, \mathbf{z}) p(\mathbf{z}|\boldsymbol{\beta}).$$

The vector \mathbf{z} cannot be sampled jointly since its entries are correlated by the Potts prior. The conditional distribution of a label $z_{j,n}$, $1 < j < J$, given all the other labels $\mathbf{z}_{\setminus(j,n)}$, is

$$z_{j,n} | \mathbf{z}_{\setminus(j,n)}, \mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Theta}, \boldsymbol{\beta} \propto p(z_{j,n} | \mathbf{z}_{\setminus(j,n)}, \boldsymbol{\beta}) p(\mathbf{x}, \boldsymbol{\mu} | \boldsymbol{\Theta}, \mathbf{z}_{\setminus(j,n)}) \quad (29)$$

with

$$p(z_{j,n} | \mathbf{z}_{\setminus(j,n)}, \boldsymbol{\beta}) \triangleq \exp \left[\sum_{m \in \mathcal{V}(n)} \beta_{xy}^j \delta(z_{j,n} - z_{j,m}) + \sum_{p \in \mathcal{D}(n)} \beta_s \delta(z_{j,n} - z_{j+1,p}) + \sum_{q \in \mathcal{U}(n)} \beta_s \delta(z_{j,n} - z_{j-1,q}) \right].$$

Using (9) and dropping all the terms independent of $z_{j,n}$, the following result is obtained:

$$p(\mathbf{x}, \boldsymbol{\mu} | \boldsymbol{\Theta}, \mathbf{z}_{\setminus(j,n)}) \propto p(\ell_{j,n} | z_{j,n}, \boldsymbol{\Theta}),$$

and the marginal density of $\ell_{j,n}$ is:

$$p(\ell_{j,n} | z_{j,n} = k, \boldsymbol{\Theta}) \sim \mathcal{N}(0, \sigma_j^2(\boldsymbol{\theta}_k)), \quad (30)$$

where $\mathcal{N}(0, \sigma^2)$ denotes a zero-mean Gaussian distribution with variance σ^2 , and $\sigma_j^2(\boldsymbol{\theta}_k) = \theta_2^k + \theta_1^k \log 2^j$.

The multiscale labels are efficiently sampled using a checker-board parallelized strategy [47, 45], here adapted to work at multiple resolutions for the multiscale label prior, as illustrated in Fig.2-(b). The proposed strategy decomposes the set of labels at a given scale j into two independent subsets, containing odd and even positions (represented in black and white in Fig.2-(b)). All labels within a subset are conditionally independent and can thus be sampled in parallel, with the labels of the other subset fixed, allowing odd and even labels to be sampled alternatively.

5.2 Sampling the Granularity Parameters β_s, β_{xy}^j

It has been shown in [45] that it is possible to approximate the conditional density $f(\boldsymbol{\beta}|\boldsymbol{\Theta}, \mathbf{x}, \boldsymbol{\mu}, \mathbf{z}) = f(\boldsymbol{\beta}|\mathbf{z})$ by considering an auxiliary vector \mathbf{w} that verifies $\boldsymbol{\Phi}(\mathbf{z}) = \boldsymbol{\Phi}(\mathbf{w})$, where $\boldsymbol{\Phi}(\mathbf{z})$ is a tractable sufficient statistic. Then, by considering the Euclidean distance $\boldsymbol{\Phi}(\mathbf{z}) - \boldsymbol{\Phi}(\mathbf{w})$, an approximation of the gradient can be computed. In practice, \mathbf{w} is initialized with \mathbf{z} , and then updated with few Gibbs moves associated with the target density $f(\mathbf{z}|\boldsymbol{\beta})$ (for more details see [46, 45]). This approximation of the gradient results in values around the neighborhood of $\boldsymbol{\beta}$ when choosing the correct statistic $\boldsymbol{\Phi}(\mathbf{z})$. Inspired by [45], we propose the statistic $\boldsymbol{\Phi}(\mathbf{z}) = \sum_{j=j_1}^{j_2} \Phi_{\beta_{xy}^j}(\mathbf{z}_j) + \Phi_{\beta_s}(\mathbf{z}_j, \mathbf{z}_{-j})$, for the scale granularity, where $\Phi_{\beta_{xy}^j}$ and Φ_{β_s} are defined in (25) and (26), and, for the spatial parameters, we use the statistic $\Phi_{\beta_{xy}^j}(\mathbf{z}_j)$, $j = j_1, \dots, j_2$. A truncation operation $T_{[0,Q]}(x)$ is introduced to restrict the parameters to belong to the interval $(0, Q)$, $Q \in \mathbb{N}$. Following this heuristic but effective strategy, the granularity parameters are sampled as summarized in Algorithm 1. Intuitively, the sampler performs gradient descent to update the $\boldsymbol{\beta}$ parameters. It first generates auxiliary variables \mathbf{w}_j using Gibbs sampling. Then, the differences between the statistics evaluated in \mathbf{z} and \mathbf{w} are computed to approximate the intractable gradient, which are used to update the values of $\boldsymbol{\beta}$.

5.3 Conditional Probability $f(\boldsymbol{\Theta}|\mathbf{z}, \boldsymbol{\beta}, \mathbf{x}, \boldsymbol{\mu})$

The conditional distributions of the multifractal parameters contained in $\boldsymbol{\Theta}$ and the latent variables in $\boldsymbol{\mu}$ are:

$$\begin{aligned} \boldsymbol{\mu}_k | \mathbf{x}_k, \theta_1^k, \theta_2^k &\sim \mathcal{CN}\left(\theta_1^k \tilde{G}_{1,k} \mathbf{x}_k, \left((\theta_1^k \tilde{G}_{1,k})^{-1} + (\theta_2^k \tilde{G}_{2,k})^{-1}\right)\right), \\ \theta_1^k | \mathbf{x}_k, \boldsymbol{\mu}_k, \theta_2^k &\sim \mathcal{IG}\left(\alpha_{1,k} + S, \gamma_{1,k} + \|\mathbf{x}_k - \boldsymbol{\mu}_k\|_{\tilde{G}_{1,k}}\right), \\ \theta_2^k | \mathbf{x}_k, \boldsymbol{\mu}_k, \theta_1^k &\sim \mathcal{IG}\left(\alpha_{2,k} + S, \gamma_{2,k} + \|\boldsymbol{\mu}_k\|_{\tilde{G}_{2,k}}\right), \end{aligned} \quad (31)$$

Algorithm 1 Multiscale granularity parameter sampling

Input: $\mathbf{z}^{(t)}, \boldsymbol{\beta}^{(t)}$: current label and granularity vector
 Q : constant (default $Q = 10$)
 t, V : current iteration, number of iterations
 Set $\mathbf{z} \leftarrow \mathbf{z}^{(t)}, [\beta_s, \beta_{xy}^{j_1}, \dots, \beta_{xy}^{j_2}] \leftarrow \boldsymbol{\beta}^{(t)}$
for $r = 1$ **to** V **do**
 $\mathbf{w}_j \leftarrow$ draw Gibbs moves of \mathbf{z}_j via (29), for $j = j_1, \dots, j_2$
 Set $\mathbf{w} = [\mathbf{w}_{j_1}^T, \dots, \mathbf{w}_{j_2}^T]^T$
 $\eta = 10(t + r - 1)^{-3/4} (\sum_{j=j_1}^{j_2} N_j)^{-1}$
 Compute
 $(\beta_s)_{(r+1)} = T_{[0, Q]} \left((\beta_s)_{(r)} + \eta(\Phi(\mathbf{z}) - \Phi(\mathbf{w})) \right)$
 for $j = j_1$ **to** j_2 **do**
 $\mathbf{b}_j = \Phi_{\beta_{xy}^j}(\mathbf{z}_j) - \Phi_{\beta_{xy}^j}(\mathbf{w}_j)$
 Compute $(\beta_{xy}^j)_{(r+1)} = T_{[0, Q]} \left((\beta_{xy}^j)_{(r)} + \eta \mathbf{b}_j \right)$
 end
end
return $\boldsymbol{\beta} = [\beta_s, \beta_{xy}^{j_1}, \dots, \beta_{xy}^{j_2}]$

where $\|\mathbf{y}\|_{\mathbf{G}} \triangleq \mathbf{y}^H \mathbf{G} \mathbf{y}$, $\tilde{G}_{i,k} = \mathbf{W}_{s,k} G_{i,s}^{-1}$.

5.4 Gibbs Sampler and Bayesian Estimators

Algorithm 2 presents the procedure for jointly sampling the labels, granularity coefficient, and multifractal parameters. The inputs of the algorithm are the image and hyperparameters assumed to be known beforehand: number of regions K , constants $\alpha_{1,k}$, $\alpha_{2,k}$, total number of generated samples N_m , number of samples of the burn-in period N_b . The output of the Gibbs sampler are used to build estimators of the parameters and labels. More precisely, the marginal posterior mean (aka minimum mean square error (MMSE)) estimator is used for the multifractal parameters $\boldsymbol{\theta}_i^k$, $i = 1, 2$:

$$(\hat{\boldsymbol{\theta}}_i^k)^{\text{MMSE}} \approx \frac{1}{(N_m - N_b)} \sum_{t=N_b+1}^{N_m} [\boldsymbol{\theta}_i^k]^{(t)}. \quad (32)$$

On the other hand, due to the discrete nature of the labels, we propose to estimate them using the maximum a-posteriori (MAP) estimator:

$$\hat{z}_{j,n}^{\text{MAP}} = \arg \max_{z_{j,n} \in \{1, \dots, K\}} \left\{ p(z_{j,n} | \mathbf{z}_{\setminus(j,n)}, \mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Theta}, \boldsymbol{\beta}) \right\}, \quad (33)$$

for $j = j_1, \dots, j_2$. Note that the computed labels are downsampled by a factor 2^{j_1} compared to the input image. In practice, the labels computed at the finest scale $\hat{\mathbf{z}}_{j_1} \in \mathbb{R}^{N_1 \times N_1}$ are interpolated (using, e.g., nearest neighbors) up to $N \times N$ to obtain labels for each pixel of the image, see [29] for alternative strategies. Finally, the updates for the granularity parameter vector $\boldsymbol{\beta}$ are performed only during the burn-in period of the sampler, as suggested in [46], in order to reduce the computational complexity. After the burn-in period, $\boldsymbol{\beta}$ is fixed to its current value.

5.5 Label Initialization

An effective approach for initializing labels consists of estimating the parameters (θ_1, θ_2) locally from the log-leaders corresponding to overlapped patches using the regression (5), where local multifractal estimates are computed on 16×16 overlapping patches (with 75% overlap). The patch-wise estimates for θ_1 are used to obtain an initial guess of the labels by k -means clustering.

Algorithm 2 Gibbs sampler

Input: \mathbf{F}, K : input image, number of regions

j_1, j_2 : initial and final scale

N_m, N_b : number of iterations, burn-in period

V : number of iterations for β (default $V = 2$)

Initialization: $\mathbf{z}^{(0)}, \Theta^{(0)}, \beta^{(0)}, \mu^{(0)}$.

Compute J log-leaders ℓ from \mathbf{F} (section 2.3).

for $t = 1$ **to** N_m **do**

 Compute Fourier coefficients $\mathbf{x} = [\mathbf{x}_1^T, \dots, \mathbf{x}_K^T]^T$ using $\mathbf{z}^{(t-1)}$ via (16).

 Sample $\mathbf{z}^{(t)} \sim p[\mathbf{z} | \Theta^{(t-1)}, \mathbf{z}^{(t-1)}, \beta^{(t-1)}, \mathbf{x}, \mu]$ via (29).

 Sample $\Theta^{(t)} \sim f(\Theta | \Theta^{(t-1)}, \mathbf{z}^{(t)}, \beta^{(t-1)}, \mathbf{x}, \mu)$, and μ via (31).

if $t < N_b$ **then**

 Generate $\beta^{(t)} \sim f(\beta | \beta^{(t-1)}, \mathbf{z}^{(t)})$ following Alg. 1.

else

$\beta^{(t)} \leftarrow \beta^{(t-1)}$.

end

end

5.6 Computational complexity

The computational complexity of the proposed approach is dominated by that of the Fourier transforms. Since $N_j \propto N$ for all j (where \propto means “proportional to”), the estimation of the multifractal parameter requires $\mathcal{O}(N^2 \log N)$ operations, performed K times. The cost of constructing the debiasing matrix in (19) is $\mathcal{O}(N^2)$. The multiscale Potts model update also requires $\mathcal{O}(N^2)$ operations to compute the probability of wavelet leaders to belong to each class. Consequently, the final complexity of the proposed method is $\mathcal{O}(KN^2 \log N)$. Although demanding for high-resolution images, computational efficiency can be improved by implementing the debiasing matrix at a lower level (e.g., C/C++) and parallelizing the multiscale Potts updates on a GPU.

6 Numerical Experiments

The proposed joint Bayesian segmentation and parameter estimation approach is first evaluated on synthetic 2D images constructed using multifractal random walks (MRWs) by means of Monte Carlo simulations¹. Some results obtained using real-world data are then presented. MRWs are non-Gaussian processes with stationary increments and are constructed in order to have multifractal properties mimicking those of Mandelbrot’s log-normal cascades, see [48] for details.

6.1 Monte Carlo Simulations

6.1.1 Synthetic Data

To validate the proposed method under controlled conditions, two scenarios with known ground-truth labels and multifractal parameters were considered. Specifically, the first scenario has $K = 2$ different multifractal regions: a background with multifractality parameter $\theta_1^1 = -0.02$ and a centered disk with parameter $\theta_1^2 = -0.08$, whereas the second scenario has $K = 3$ regions: a background with $\theta_1^1 = -0.02$ and two disks with parameters $\theta_1^2 = -0.08$ and $\theta_1^3 = -0.16$. The experiments were run on 100 independent realizations of MRWs with these parameters, where the computation time for one realization is in average 58 seconds². An illustration of a single realization of an image for each scenario and their corresponding true masks is displayed in Figs. 4 (a) and (b), respectively.

¹The source code is available at https://github.com/kareth-leon/MFA_Seg

²The simulations were conducted in a desktop architecture with an Intel Corei7, 16 GB RAM, using MATLAB R2024b.

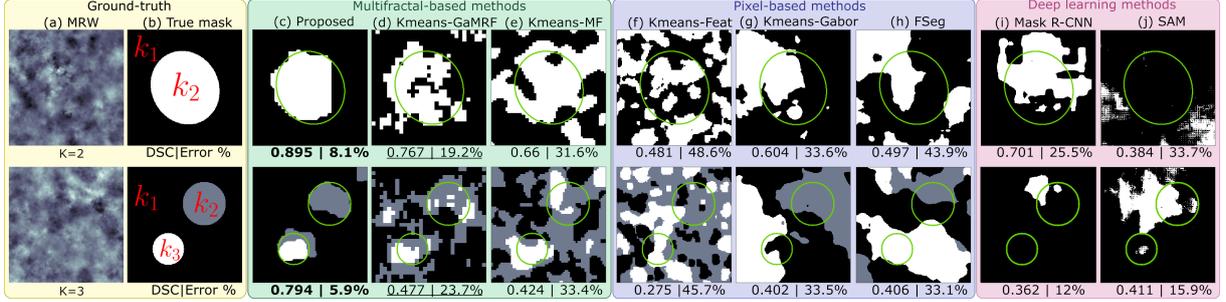


Figure 4: Single MRW realization for $K = 2$ (top row) and $K = 3$ (bottom row) and their segmentation labels obtained from the different compared methods: (a) MRW image generated using the true labels in (b). The proposed segmentation (c) is compared against methods using multifractal parameters (d) and (e), framed in the green square; classical pixel-based segmentation methods: (f) KMEANS-FEAT, (g) KMEANS-GABOR, (h) FSEG, framed in the blue square; and deep-based segmentation approaches: (i) Mask R-CNN, and (j) SAM, framed in the rose square. The average DSC and error percentage of each results are shown in the bottom, where the scores obtained by the proposed method (in bold) are superior regarding the others.

6.1.2 Performance Measure

The performance of the multifractality parameter estimation is quantified using the mean, the standard deviation (STD) and the root-mean square error (RMSE) of the estimates. The label estimation performance is assessed using the dice score coefficient (DSC) and the segmentation error (aka misclassification error), which are defined in terms of true positive (TP), true negative (TN), false positive (FP), and false negative (FN): DSC (aka F1-score) is defined on each region as $DSC = 2 TP / (2 TP + FP + FN)$. The higher the DSC score, the better the segmentation result. The segmentation error is defined by $ERROR = (FP+FN)/(TP+TN+FP+FN)$. The lower the error, the better the performance.

6.1.3 Parameter Settings

The Daubechies mother wavelet with $N_\psi = 1$ vanishing moments is used in the 2D DWT. The hyperparameters of the independent priors for the multifractal parameters were set to $(\alpha_{i,k}, \gamma_{i,k}) = (10^{-3}, 10^{-3})$ [4]. In practice, these constants are fixed to small values, ensuring that they are close to non-informative Jeffreys priors [4, 45]. The set of scales used for the experiment with $K = 2$ classes was set to $(j_1, j_2) = (1, 3)$, whereas $(j_1, j_2) = (1, 2)$ for $K = 3$. Note that j_2 is limited by image dimensions ($j_2 < \log_2 N$), nonetheless, in this study, j_2 was set between 2 and 3, as the coarsest scales were found to retain meaningful pixel-level information. The number of iterations for all the experiments was set to $N_m = 300$, with a burn-in period of $N_b = 30$, since it showed to be sufficient for sampling convergence (see next subsection for convergence details).

6.2 Estimation Performance

6.2.1 Multifractal Parameter Estimation on Irregular Shapes

The proposed approach is first evaluated in a setting where the labels are known and fixed with $K = 2$ classes, in order to assess the estimation performance of the multifractal parameter on an irregular domain. Furthermore, to illustrate the impact of the multifractality parameter value on the image segmentation, the multifractality value in the disk of the MRWs was varied in a range of weak (close to zero) to strong values $\theta_1^2 \in \{-0.005, -0.08, -0.2\}$, fixing the background parameter to $\theta_1^1 = -0.02$. Results are reported in Table 1. They indicate that the proposed approach yields satisfactory performance (small RMSE values), comparable to that obtained for regular shapes [4, 27].

6.2.2 Multifractal Parameter and Label Estimation

Table 2 shows the performance of the multifractal parameter estimation when the labels and granularity parameters are also unknown and estimated jointly, following the previously described setting. The results

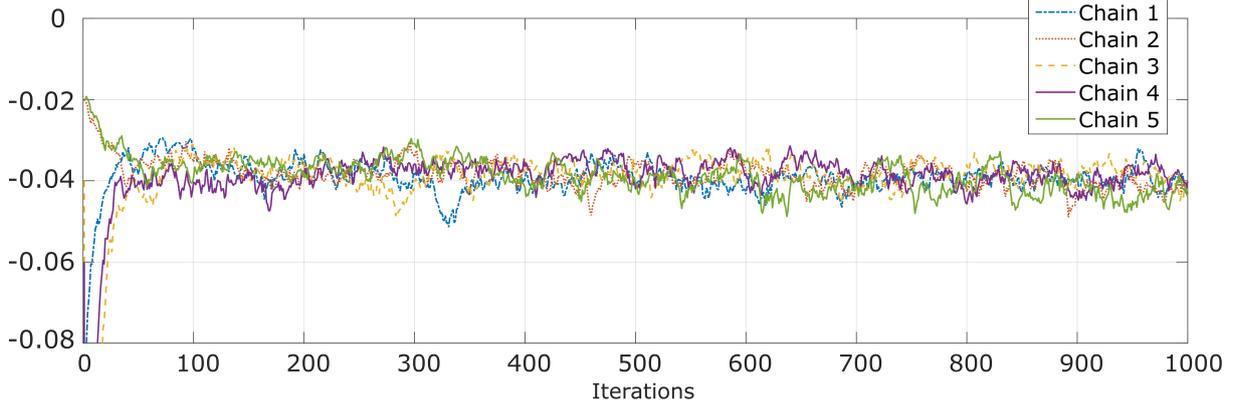


Figure 5: Convergence assessment with five Markov chains with different initializations.

Table 1: Performance of the multifractal parameter estimation $(\tilde{\theta}_1^k)^{\text{MMSE}}$ when the labels are fixed and known.

k	True θ_1^k	MEAN (STD)	RMSE
1	-0.02	-0.032 (0.005)	0.013
2	-0.005	-0.014 (0.003)	0.010
1	-0.02	-0.034 (0.005)	0.015
2	-0.08	-0.087 (0.016)	0.018
1	-0.02	-0.034 (0.004)	0.015
2	-0.2	-0.184 (0.028)	0.032

show that the RMSE values slightly increase, due to the larger number of unknowns to be estimated, but that the overall performance remains comparable to Table 1. Table 3 reports the DSC values and the segmentation errors for the two experiments, allowing the label estimation performance to be evaluated. The results indicate that the segmentation error decreases as the difference in the multifractal parameter between the regions increases. This behavior is indeed expected, since the increased parameter contrast should facilitate segmentation. Overall, the segmentation performance is very satisfactory, except for the extreme case $\theta_1^2 = -0.005$ that corresponds to nearly no multifractality.

6.2.3 Sampler Convergence

To evaluate the convergence of the proposed method, we use the potential scale reduction factor (PSRF), a convergence diagnostic for Monte Carlo simulations [49]. It compares the variance within the chain to the variance between chains in multiple chains, where $\text{PSRF} < 1.2$ indicates good convergence of the sampler (see [49, 50] for further details). Any parameter sampled in the Monte Carlo simulations can serve as a reference to monitor convergence. For illustration purposes, we chose to monitor the multifractal value of the background of Fig. 4-a, $\theta_1^1 = -0.02$, during 1000 iterations. Figure 5 shows five chains of the parameter with different initial values. It can be seen that the chains converge to similar values, showing that the method converges even when changing the initial value of the chain. The PSRF value obtained for this parameter with 1000 Monte Carlo runs is 1.0009, exhibiting the convergence of the sampler.

6.3 Comparison to the State of the Art

The proposed segmentation approach is compared to unsupervised state-of-the-art methods for texture/image segmentation, including deep learning methods, classical texture segmentation, and the classical k -means clustering method, a benchmark in segmentation, used to estimate labels for different features, including multifractal, pixel intensity, and multiscale features. These methods are summarized below:

- The k -means on multifractal parameters estimated by regression (KMEANS-MF). Note that this method was used to initialize the suggested method described in 5.5.

Table 2: Performance of the multifractal parameter estimation $(\tilde{\theta}_1^k)^{\text{MMSE}}$ when labels and parameters are jointly estimated.

k	True θ_1^k	MEAN (STD)	RMSE
1	-0.02	-0.041 (0.014)	0.025
2	-0.005	-0.023 (0.013)	0.022
1	-0.02	-0.033 (0.006)	0.014
2	-0.08	-0.113 (0.025)	0.041
1	-0.02	-0.037 (0.011)	0.020
2	-0.2	-0.238 (0.058)	0.069

Table 3: Segmentation performance for the labels in terms of DSC and Segmentation Error in the joint parameters and label estimation scenario for 100 realizations by considering different multifractal parameter $\theta_1^{(k=2)}$ in the disk.

$[\theta_1^{(k=1)}, \theta_1^{(k=2)}]$	DSC \uparrow [$k = 1$]	DSC \uparrow [$k = 2$]	ERROR \downarrow %
$[-0.02, -0.005]$	0.649 (0.103)	0.487 (0.223)	39.4 (8.3)
$[-0.02, -0.08]$	0.909 (0.034)	0.752 (0.117)	13.2 (4.8)
$[-0.02, -0.2]$	0.940 (0.030)	0.816 (0.140)	9.0 (5.0)

- The k -means on a localized Bayesian multifractal estimates (KMEANS-GAMRF) [4], which makes use of a smoothing gamma Markov random field (GaMRF) prior³. The resulting multifractal estimates are used as features for k -means segmentation.
- The k -means on standard patch-wise image features (KMEANS-FEAT): uses patch-wise features including the mean, standard deviation, minimum and maximum values of patches of size 10×10 with 80% overlap.
- The k -means on Gabor filter features (KMEANS-GABOR) [11]: a benchmark for classical multiscale features extraction without multifractal model.
- The factorization-based texture segmentation (FSEG) method [16]: computes multiscale Gabor coefficients followed by a feature selection procedure by a matrix factorization step. This method is selected given its multiscale nature for segmentation. The hyperparameters were adjusted as in [16].
- The segment anything (SAM) method [12]: it relies on a *promptable* segmentation that takes as inputs the image and a handcrafted reference (such as points, boxes or mask logits), and returns a segmentation mask [12]. The pre-trained SAM model was supervised by providing the spatial coordinates of two pixels randomly selected from each region.
- The MASK R-CNN architecture [13]: a pre-trained convolutional neural network for instance segmentation that generates a pixel-level binary mask for each detected object in the given image. A threshold of 0.98 was applied to the network’s output masks for the detected objects.

The SAM and MASK R-CNN models were selected since they are outstanding methods of the state of the art in deep-based image/instance segmentation. Additionally, note that the SAM model relies on multiple transformers and attention modules that model long-range dependencies to handle global context and complex scenes [12].

Figure 4 displays the segmentation results obtained using the different methods for a single MRW realization, when $K = 2$ (top row) and $K = 3$ (bottom row). The proposed method yields an average DSC close to 0.85, significantly outperforming the other methods, and provides visually satisfactory results. All other baselines lead to significantly worse results, illustrating the difficulty of segmenting multifractal images. Tables 4 and 5 provide a more quantitative and systematic comparison in terms of average performance for 100 independent realizations of MRW for $K = 2$ and 3. It can be seen that the proposed method significantly outperforms the other methods, showing overall best segmentation scores in terms of DSC and ERROR by a large margin.

³Toolbox available at <https://www.irit.fr/~Herwig.Wendt/CIMPA2017.html>

Table 4: Segmentation baseline comparison for the label estimation when $K = 2$ in terms of average DSC and segmentation Error percentage of 100 realizations, the standard deviation is shown in parenthesis. The best scores are highlighted in bold.

Method	DSC \uparrow [$k = 1$]	DSC \uparrow [$k = 2$]	ERROR \downarrow %
PROPOSED	0.909 (0.034)	0.752 (0.117)	13.2 (4.8)
KMEANS-GAMRF	0.843 (0.027)	0.567 (0.097)	22.8 (3.7)
KMEANS-MF	0.745 (0.079)	0.539 (0.111)	32.2 (7.9)
KMEANS-FEAT	0.643 (0.071)	0.383 (0.088)	44.2 (5.7)
KMEANS-GABOR	0.683 (0.084)	0.443 (0.126)	39.6 (8.4)
FSEG	0.637 (0.050)	0.416 (0.077)	44.6 (5.2)
MASK-RCNN	0.813 (0.042)	0.300 (0.250)	28.5 (0.1)
SAM	0.196 (0.215)	0.458 (0.054)	62.7 (8.8)

Table 5: Average DSC for 100 realizations of the baseline comparison when $K = 3$. The standard deviation is shown in parenthesis.

	DSC \uparrow [$k = 1$]	DSC \uparrow [$k = 2$]	DSC \uparrow [$k = 3$]
PROPOSED	0.846 (0.072)	0.408 (0.147)	0.494 (0.256)
KMEANS-GAMRF	0.846 (0.051)	0.328 (0.124)	0.439 (0.173)
KMEANS-MF	0.669 (0.094)	0.230 (0.083)	0.390 (0.147)
KMEANS-FEAT	0.505 (0.150)	0.219 (0.068)	0.139 (0.078)
KMEANS-GABOR	0.515 (0.091)	0.353 (0.102)	0.210 (0.119)
FSEG	0.557 (0.118)	0.322 (0.113)	0.198 (0.141)
MASK-RCNN	0.765 (0.292)	0.093 (0.166)	0.037 (0.112)
SAM	0.547 (0.186)	0.282 (0.112)	0.016 (0.088)

6.4 Experiments using real data

This section studies the effectiveness of the proposed approach for real data segmentation. The proposed segmentation method is first applied to a mammogram image for tumor segmentation in breast screening, and second, to a synthetic aperture radar (SAR) image for water body detection.

Mammography remains the primary technique for detecting and diagnosing breast cancer. Early cancer detection can significantly reduce treatment costs. Therefore, the development of methodologies that enable accurate tumor detection is of great interest, with promising results obtained using multifractal-based methods [23]. On the other hand, water body segmentation based on SAR images has become a research hotspot in remote sensing applications given that SAR satellites are able to operate any time, regardless of weather or lighting conditions [51, 52]. However, SAR images are contaminated by multiplicative speckle noise, making their use difficult for computer vision tasks [52]. From these scenarios, it is important to remark that the applicability of the proposed method can be determined by assessing whether the data exhibit multifractal, scale-free behavior. This can be verified by checking if the scaling relation (3) [27] holds empirically, which is the case for the selected scenarios.

6.4.1 Tumor segmentation

The mammography image used for this experiment was extracted from the Mini-DDSM (Digital Database of Screening Mammography) dataset⁴. In this dataset, each image has a segmentation label, or contour, that delineates the margins of the tumor⁵. In this context, the proposed method can be considered with

⁴Mammography dataset available at <https://ardisdataset.github.io/MiniDDSM/>. The set used for the simulation is “cancer/1727”.

⁵It is important to note that the provided contour serves just as a reference to localize the tumor, as literature has reported that such contours may exhibit inaccuracies. Nevertheless, they offer a sufficiently accurate approximation for characterizing the tumor. The segmentation mask is computed from the provided contour by

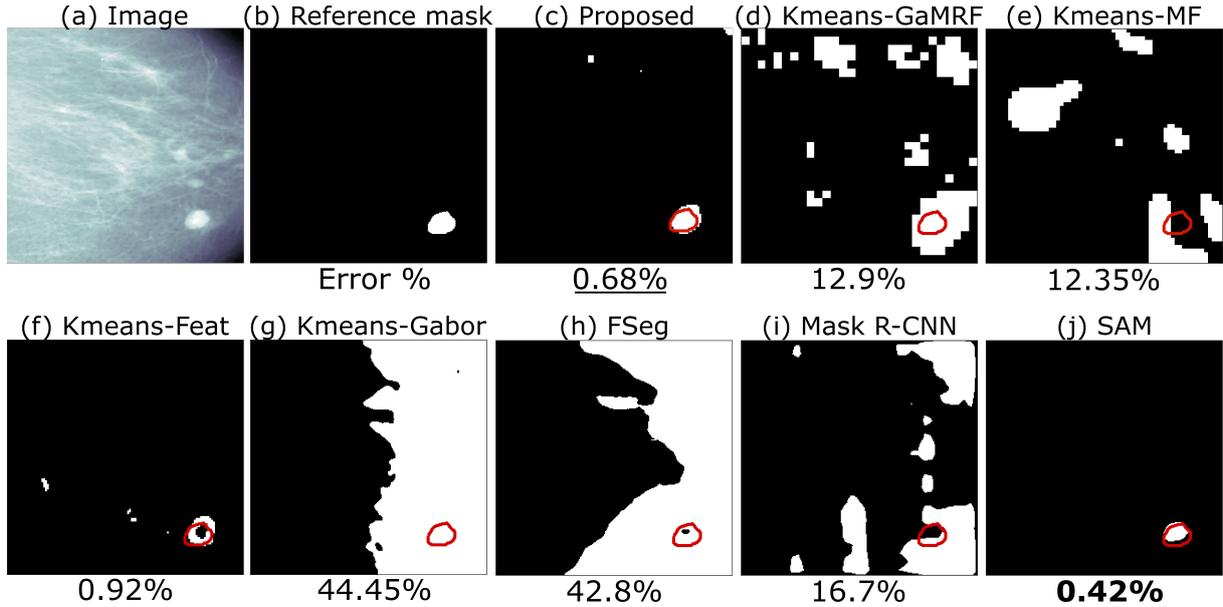


Figure 6: Comparison of segmentation results on a mammography image for tumor detection, with the tumor contour shown in red on the resulting masks and the segmentation error displayed below each result. Notably, SAM method achieves the best segmentation performance, yet it requires ground-truth pixels from each region as prompts for its operation, making it supervised. The second best segmentation is achieved with the proposed method, which is fully unsupervised, showing its capability of modeling different multifractalities on an image for segmentation purposes.

$K = 2$ assuming that tumor and normal tissue are characterized by different multifractal parameters. The mammography image was cropped to retain only the region of interest shown in Figure 6-(a), removing the surrounding background. Figure 6 shows the analyzed section of (a) the mammography image with (b) its respective segmentation mask, highlighting the tumor, and (c-j) the segmentation results of the proposed method compared to the baseline methods. The proposed method yields very satisfactory results, with the second best error (0.68%) after the SAM method (0.42%). Note that SAM produces more false negatives in the region of interest (FN: 203) than the proposed method (FN: 53). Also note that SAM operates in a supervised manner and requires knowledge of input pixels from each region. In contrast, the proposed method operates in a fully unsupervised mode, demonstrating its potential effectiveness in modeling and handling tumor segmentation tasks.

6.4.2 Water detection from SAR images

Two SAR images were extracted from the dataset in [51], where the true label for the water segmentation is provided (we considered $K = 2$). Figure 7 presents a comparison of segmentation performance, showing that the proposed unsupervised method (c) achieves competitive results and is second best after the FSEG (h) and SAM methods (j): SAM, which is supervised (two points per class were provided as input), yields better performance for the second image but significantly worse for the first, where it fails to detect one of the two water bodies. The FSEG method shows outstanding performance for the first image, but fails quite drastically for the second one. In contrast, the proposed method well detects the water bodies in all cases and yields very satisfactory and robust performance. These results underline the robustness and versatility of the proposed method to handle different segmentation scenarios on challenging images such as SAR data.

performing a flood-fill operation on the binary mask.

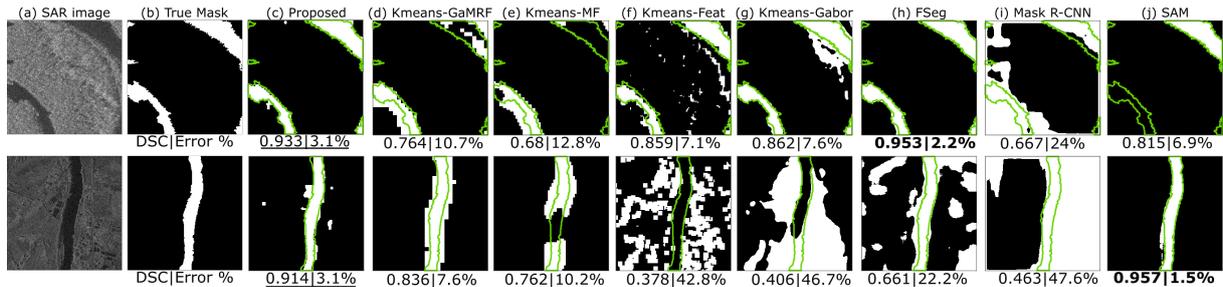


Figure 7: Results on SAR images for segmenting water bodies. First and second best performances are highlighted in bold and with an underline, respectively. For SAM, two points per class were provided as prompts to guide the segmentation. All the other methods are unsupervised.

7 Conclusions

This work introduced a new Bayesian multifractal image segmentation method formulated as the joint estimation of the multifractality parameters and pixel labels. As a novel key ingredient, the proposed approach leverages on a flexible Fourier domain model for log-leaders allowing the multifractality estimation of heterogeneous images. To model the space-scale label correlations, a multiscale graph structure was introduced and used in a Potts Markov random field prior. A Bayesian formulation was finally introduced to estimate the unknown model parameters (including the multifractality of each region of the image, the image labels and the granularity parameters of the image regions), numerically and efficiently computed using samples generated by an appropriate Gibbs sampler. Numerical experiments conducted on synthetic multifractal images and real data demonstrate the excellent performance of the proposed approach for image segmentation, outperforming several baselines including local multifractality-based approaches, classical texture segmentation methods as well as recent deep learning strategies. Future work includes reducing computational time through parallel implementations (e.g., GPUs) and the extension of the model to data with more than two dimensions (such as depth or spectral data volumes [53]), in which multidimensional/multivariate correlations can be included and exploited in the model. Other potential directions are the evaluation of the proposed approach in scenarios with corrupted spatial information, such as noisy, missing or compressed data (as in [54]), and the integration of the proposed method with deep learning networks (as in [55]).

References

- [1] P. Abry, H. Wendt, and S. Jaffard, “When Van gogh meets Mandelbrot: multifractal classification of painting’s texture,” *Signal Processing*, vol. 93, no. 3, pp. 554–572, 2013.
- [2] H. Wendt, P. Abry, and S. Jaffard, “Bootstrap for empirical multifractal analysis,” *IEEE Trans. Signal Proces. Mag.*, vol. 24, no. 4, pp. 38–48, 2007.
- [3] H. Wendt and P. Abry, “Multifractality tests using bootstrapped wavelet leaders,” *IEEE Trans. Signal Proces.*, vol. 55, no. 10, pp. 4811–4820, 2007.
- [4] H. Wendt, S. Combrexelle, Y. Altmann, J.-Y. Tourneret, S. McLaughlin, and P. Abry, “Multifractal analysis of multivariate images using gamma Markov random field priors,” *SIAM Imaging Sciences*, vol. 11, no. 2, pp. 1294–1316, 2018.
- [5] P. Abry and D. Veitch, “Wavelet analysis of long-range-dependent traffic,” *IEEE Trans. Info. Theory*, vol. 44, no. 1, pp. 2–15, Jan. 1998.
- [6] H. Wendt, P. Abry, S. Jaffard, H. Ji, and Z. Shen, “Wavelet leader multifractal analysis for texture classification,” in *Proc. IEEE Int. Conf. Image Proces. (ICIP)*, Cairo, Egypt, Nov. 2009, pp. 3829–3832.
- [7] H. Wendt, S. G. Roux, S. Jaffard, and P. Abry, “Wavelet leaders and bootstrap for multifractal analysis of images,” *Signal Proces.*, vol. 89, no. 6, pp. 1100–1114, 2009.

- [8] Y. Xu, X. Yang, H. Ling, and H. Ji, “A new texture descriptor using multifractal analysis in multi-orientation wavelet pyramid,” in *Proc. IEEE Conf. Comp. Vis. Pattern Recognition (CVPR)*, San Francisco, CA, USA, June 2010, pp. 161–168.
- [9] F. G. Schmitt and L. Seuront, “Multifractal random walk in copepod behavior,” *Physica A: Statistical Mechanics and its Applications*, vol. 301, no. 1-4, pp. 375–396, 2001.
- [10] S. Jaffard, “Wavelet techniques in multifractal analysis,” in *Fractal Geometry and Applications: A Jubilee of Benoit Mandelbrot*, M. Lapidus and M. van Frankenhuysen, Eds. in *Proc. Symposia Pure Mathematics*, vol. 72, no. 2. American Mathematical Society, 2004, pp. 91–152.
- [11] A. K. Jain and F. Farrokhnia, “Unsupervised texture segmentation using Gabor filters,” *Pattern Recognition*, vol. 24, no. 12, pp. 1167–1186, 1991.
- [12] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo *et al.*, “Segment anything,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Paris, France, Oct. 2023, pp. 4015–4026.
- [13] K. He, G. Gkioxari, P. Dollár, and R. Girshick, “Mask R-CNN,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Venice, Italy, Oct. 2017, pp. 2961–2969.
- [14] B. Manjunath and R. Chellappa, “Unsupervised texture segmentation using Markov random field models,” *IEEE Trans. Pattern Analysis Machine Intelligence*, vol. 13, no. 5, pp. 478–482, 1991.
- [15] A. Halimi, N. Dobigeon, and J.-Y. Tourneret, “Unsupervised unmixing of hyperspectral images accounting for endmember variability,” *IEEE Trans. Image Proces.*, vol. 24, no. 12, pp. 4904–4917, 2015.
- [16] J. Yuan, D. Wang, and A. M. Cheriyyadat, “Factorization-based texture segmentation,” *IEEE Trans. Image Proces.*, vol. 24, no. 11, pp. 3488–3497, 2015.
- [17] R. M. Haralick, “Statistical and structural approaches to texture,” *Proceedings of the IEEE*, vol. 67, no. 5, pp. 786–804, 1979.
- [18] N. Pustelnik, H. Wendt, P. Abry, and N. Dobigeon, “Combining local regularity estimation and total variation optimization for scale-free texture segmentation,” *IEEE Trans. Comput. Imaging*, vol. 2, no. 4, pp. 468–479, 2016.
- [19] M. Abadi and E. Grandchamp, “Texture features and segmentation based on multifractal approach,” in *Prog. in Patt. Rec., Im. Ana. and App.: 11th Iber. Cong. in Patt. Recog.* Cancun, Mexico: Springer, Nov. 14-17 2006, pp. 297–305.
- [20] B. B. Chaudhuri and N. Sarkar, “Texture segmentation using fractal dimension,” *IEEE Trans. Pattern Anal. and Mach. Intel.*, vol. 17, no. 1, pp. 72–77, 1995.
- [21] V. M. San Martin and A. Figliola, “Application of multifractal analysis to segmentation of water bodies in optical satellite images,” *Papers in physics*, vol. 9, no. 2, pp. 1–10, 2017.
- [22] L. Zhu, G. Xiong, D. Guo, and W. Yu, “Ship target detection and segmentation method based on multi-fractal analysis,” *The Journal of Eng.*, vol. 2019, no. 21, pp. 7876–7879, 2019.
- [23] R. Lopes and N. Betrouni, “Fractal and multifractal analysis: A review,” *Medical Image Analysis*, vol. 13, no. 4, pp. 634–649, 2009.
- [24] Y. Xia, D. Feng, and R. Zhao, “Morphology-based multifractal estimation for texture segmentation,” *IEEE Trans. on Image Proces.*, vol. 15, no. 3, pp. 614–623, 2006.
- [25] J. L. Vehe and P. Mignot, “Multifractal segmentation of images,” *Fractals*, vol. 2, no. 03, pp. 371–377, 1994.
- [26] S. Jaffard, B. Lashermes, and P. Abry, “Wavelet leaders in multifractal analysis,” in *Wavelet Anal. and App.* Springer, 2006, pp. 201–246.
- [27] S. Combrexelle, H. Wendt, N. Dobigeon, J.-Y. Tourneret, S. McLaughlin, and P. Abry, “Bayesian estimation of the multifractality parameter for image texture using a Whittle approximation,” *IEEE Trans. Image Proces.*, vol. 24, no. 8, pp. 2540–2551, 2015.

- [28] H. Wendt, L. Leon, J.-Y. Tournieret, and P. Abry, “Multifractal anomaly detection in images via space-scale surrogates,” in *Proc. IEEE Int. Conf. Image Proces. (ICIP)*, Bordeaux, France, 16-19 Oct. 2022, pp. 556–560.
- [29] H. Choi and R. G. Baraniuk, “Multiscale image segmentation using wavelet-domain hidden Markov models,” *IEEE Trans. Image Proces.*, vol. 10, no. 9, pp. 1309–1321, 2001.
- [30] P. C. Brault and A. Mohammad-Djafari, “Unsupervised Bayesian wavelet domain segmentation using Potts-Markov random field modeling,” *Journal of Electronic Imaging*, vol. 14, no. 4, p. 043011, 2005.
- [31] C. A. Bouman and M. Shapiro, “A multiscale random field model for Bayesian image segmentation,” *IEEE Trans. Image Proces.*, vol. 3, no. 2, pp. 162–177, Mar. 1994.
- [32] G. Fan and X.-G. Xia, “A joint multicontext and multiscale approach to Bayesian image segmentation,” *IEEE Trans. Geosc. Remote Sensing*, vol. 39, no. 12, pp. 2680–2688, 2001.
- [33] H. Cheng and C. A. Bouman, “Multiscale Bayesian segmentation using a trainable context model,” *IEEE Trans. Image Proces.*, vol. 10, no. 4, pp. 511–525, 2001.
- [34] T. H. Kim, I. K. Eom, and Y. S. Kim, “Multiscale bayesian texture segmentation using neural networks and Markov random fields,” *Neural Computing and Applications*, vol. 18, no. 2, pp. 141–155, Jan. 2009.
- [35] B. Pesquet-Popescu and J. L. V  hel, “Stochastic fractal models for image processing,” *IEEE Signal Proces. Mag.*, vol. 19, no. 5, pp. 48–62, 2002.
- [36] A. K. Lerner and F. Nazarov, “Intuitive dyadic calculus: the basics,” *Expositiones Mathematicae*, vol. 37, no. 3, pp. 225–265, Sept. 2019.
- [37] S. Jaffard, P. Abry, and H. Wendt, *Irregularities and Scaling in Signal and Image Processing: Multifractal Analysis*. Singapore: World scientific publishing, 2015, pp. 31–116.
- [38] B. B. Mandelbrot and J. W. Van Ness, “Fractional Brownian motions, fractional noises and applications,” *SIAM review*, vol. 10, no. 4, pp. 422–437, 1968.
- [39] P. Abry, P. Flandrin, M. S. Taqqu, and D. Veitch, *Wavelets for the analysis, estimation, and synthesis of scaling data*. Wiley Online Library, 2000, ch. 2, pp. 39–88.
- [40] L. Leon, H. Wendt, J.-Y. Tournieret, and P. Abry, “A Bayesian framework for multivariate multifractal analysis,” *IEEE Trans. Signal Proces.*, vol. 70, pp. 3663–3675, 2022.
- [41] M. Fuentes, “Approximate likelihood for large irregularly spaced spatial data,” *American Statistical Assoc.*, vol. 102, no. 477, pp. 321–331, 2007.
- [42] A. P. Guillaumin, A. M. Sykulski, S. C. Olhede, and F. J. Simons, “The debiased spatial Whittle likelihood,” *Royal Statistical Society: Series B*, vol. 84, no. 4, pp. 1526–1557, 2022.
- [43] F.-Y. Wu, “The Potts model,” *Rev. Mod. Phys.*, vol. 54, no. 1, p. 235, 1982.
- [44] M. Pereyra, N. Dobigeon, H. Batatia, and J.-Y. Tournieret, “Segmentation of skin lesions in 2-D and 3-D ultrasound images using a spatially coherent generalized rayleigh mixture model,” *IEEE Trans. Medical Imaging*, vol. 31, no. 8, pp. 1509–1520, Aug. 2012.
- [45] —, “Estimating the granularity coefficient of a Potts-Markov random field within a Markov chain Monte Carlo algorithm,” *IEEE Trans. Image Proces.*, vol. 22, no. 6, pp. 2385–2397, 2013.
- [46] M. Pereyra, N. Whiteley, C. Andrieu, and J.-Y. Tournieret, “Maximum marginal likelihood estimation of the granularity coefficient of a Potts-Markov random field within an MCMC algorithm,” in *Proc. IEEE Workshop Statistical Signal Proces. (SSP)*, Gold Coast, QLD, Australia, 29 June-02 July 2014, pp. 121–124.
- [47] J. Gonzalez, Y. Low, A. Gretton, and C. Guestrin, “Parallel Gibbs sampling: From colored fields to thin junction trees,” in *Proc. Art. Int. and Stat.*, Fort Lauderdale, FL, USA, 11–13 Apr. 2011, pp. 324–332.
- [48] R. Robert and V. Vargas, “Gaussian multiplicative chaos revisited,” *The Annals of Probability*, vol. 38, no. 2, pp. 605 – 631, 2010.
- [49] A. Gelman and D. B. Rubin, “Inference from iterative simulation using multiple sequences,” *Stat. scie.*, vol. 7, no. 4, pp. 457–472, 1992.

- [50] N. Dobigeon, J.-Y. Tournet, and J. D. Scargle, “Joint segmentation of multivariate astronomical time series: Bayesian sampling with a hierarchical model,” *IEEE Trans. on Signal Proces.*, vol. 55, no. 2, pp. 414–423, 2007.
- [51] M. U. Kim, H. Oh, S.-J. Lee, Y. Choi, and S. Han, “A large-scale dataset for water segmentation of SAR satellite,” in *IEEE/RSJ Int. Conf. on Intel. Rob. and Sys. (IROS)*. Prague, Czech Republic: IEEE, 2021, pp. 9796–9801.
- [52] L. Chen, X. Cai, J. Xing, Z. Li, W. Zhu, Z. Yuan, and Z. Fang, “Towards transparent deep learning for surface water detection from SAR imagery,” *Int. J. of App. Earth Obs. and Geoinf.*, vol. 118, p. 103287, 2023.
- [53] M. A. A. Belmekki, J. Leach, R. Tobin, G. S. Buller, S. McLaughlin, and A. Halimi, “3D target detection and spectral classification for single-photon LiDAR data,” *Optics Express*, vol. 31, no. 15, pp. 23 729–23 745, 2023.
- [54] K. M. León-López and H. Arguello, “Online tensor sparsifying transform based on temporal superpixels from compressive spectral video measurements,” *IEEE Trans. Image Proces.*, vol. 29, pp. 5953–5963, 2020.
- [55] G. Xiong, T. Zhen, W. Huang, B. Min, and W. Yu, “Fractal-domain deep learning with transformer architecture for SAR ship classification,” *ISPRS J. of Photo. and Remote Sensing*, vol. 230, pp. 208–226, 2025.