Bayesian Unsupervised Multifractal Image Segmentation using a Multiscale Graph Label Prior

Kareth León, Jean-Yves Tourneret
ENSEEIHT/IRIT/INP-Toulouse, TéSA
Toulouse, France
jean-yves.tourneret@enseeiht.fr

Abderrahim Halimi
Heriot-Watt University
Edinburgh EH14 4AS, UK
a.halimi@hw.ac.uk

Herwig Wendt
CNRS, INP-Toulouse, University of Toulouse
Toulouse, France
herwig.wendt@toulouse-inp.fr

***Abstract*—This paper presents a Bayesian multifractal segmentation method that segments multifractal textures in regions with different multifractal properties. First, a computationally and statistically efficient model for wavelet leader-based multifractal parameter estimation is developed, assigning wavelet leader coefficients associated with distinct parameters to different image regions. Next, a multiscale graph label prior is introduced to capture spatial and scale correlations among these labels. Gibbs sampling is used to generate samples from the posterior distribution. Numerical experiments on synthetic multifractal images demonstrate the effectiveness of the proposed method, outperforming traditional unsupervised and modern deep learning-based segmentation approaches.**

***Index Terms*—Texture segmentation, wavelet leaders, multifractal analysis, Bayesian estimation, Gibbs sampling.**

I. Introduction

Multifractality of Images. Multifractal analysis (MFA) has emerged as a powerful tool to characterize complex textures in images, with applications in, e.g., remote sensing [1], medial imaging [2], art investigation [3], and texture analysis [4]. MFA enables the study of *local regularity* of images or functions. Precisely, the texture of an image $\mathbf{F}$ is encoded by the *multifractal spectrum* $\mathcal{D}(h)$, which is defined as the collection of fractal dimensions (Hausdorff dimension) of the sets of points that take the same pointwise regularity value h, typically measured with the *Hölder exponent* [5], [6]. The primary objective of MFA is to estimate $\mathcal{D}(h)$, which in practice is achieved through a multiscale decomposition and the *multifractal formalism* [1], [5]. A state-of-the-art approach to MFA relies on wavelet leaders $L_{j,\mathbf{n}}$, derived from the two-dimensional (2D) discrete wavelet transform (DWT) coefficients of the image, as defined in Section II.

Estimation of the multifractality parameter c_2. It has been shown that the multifractal formalism provides an expansion of $\mathcal{D}(h)$ in terms of the so-named log-cumulants c_p, $p \geq 1$,

$$\mathcal{D}(h) = 2 + \frac{c_2}{2!}\Big(\frac{h-c_1}{c_2}\Big)^2 + \frac{-c_3}{3!}\Big(\frac{h-c_2}{c_2}\Big)^3 + ...,$$

for $c_2 < 0$ [1], [6]. The coefficients c_p are tied to the wavelet leaders $L_{j,\mathbf{n}}$, via the relation [6]

$$\mathrm{Cum}_p[\ln L_{j,\mathbf{n}}] = c_p^0 + c_p \ln 2^j, \tag{1}$$

Work partially supported by Grant ANR-18-CE45-0007 MUTATION and the UK Royal Academy of Engineering (RF/201718/17128).

and offer a concise representation of $\mathcal{D}(h)$. For instance, c_1 corresponds to the mode of $\mathcal{D}(h)$ and is related to the Hurst parameter [7], reflecting the average smoothness of $\mathbf{F}$. Meanwhile, c_2 quantifies the width of $\mathcal{D}(h)$, characterizing the extent of local regularity fluctuations in $\mathbf{F}$. Thus, c_2 is one of the key quantities of interest in MFA, referred to as the *multifractality parameter*. The relation (1) with $p = 2$ yields $\mathrm{Var}[\ln L_{j,\mathbf{n}}] = c_2^0 + c_2 \ln 2^j$, indicating that c_2 can be estimated by linear regression over a range of scales $j = j_1, ..., j_2$, $\hat{c}_2 = \frac{1}{\ln 2}\sum_{j=j_1}^{j_2} w_j \widehat{\mathrm{Var}}[\ln(L_{j,\mathbf{n}})]$, where $\widehat{\mathrm{Var}}(\cdot)$ is the sample variance and w_j are suitable regression weights [5], [8], [9]. This estimator is computationally efficient, but exhibits limited accuracy. To address this issue, an efficient Fourier-based Bayesian model and estimators was proposed for homogeneous images [1]. However, the approach is not suitable for irregularly shaped domains, hence for segmentation purposes, and can at best operate on small patches of homogeneous textures, limiting its use for segmentation purposes.

Several methods have shown the interest in using multiresolution and wavelet coefficients for image segmentation [10]–[14]. Others, such as [15], [16], have used multifractal-based features to perform segmentation. However, while these methods effectively integrate multiscale information, they typically rely on training data and ignore the multifractality of textures, driven by c_2.

Contributions. The key difficulty in the segmentation for multifractal (MF) models arises from MF parameter estimation requiring the analysis of several scales, thus being tied to space averages, while segmentation has to be performed at the pixel level. This work addresses this difficulty with a novel unsupervised MF image segmentation method based on a Bayesian framework and multiscale model. As a first key contribution, a Fourier-based likelihood in the spirit of [1] is defined for images with several MF textures (Sec. III). Second, an original Potts-Markov random field with a multiscale graph prior is proposed, capturing both spatial and scale dependencies in the labels. This enables defining a Bayesian model and estimators, approximated using a Gibbs sampler, for joint MF parameter estimation and label-based segmentation (Sec. IV). Numerical experiments are conducted for synthetic MF images and demonstrate the good performance of the proposed approach over state-of-the-art segmentation methods (Sec. V).

II. Fourier-based Model for the Estimation of c_2

Wavelet Leaders. Let $D_{\mathbf{F}}^{(m)}(j,\mathbf{n})$ denote the 2D DWT of the image $\mathbf{F} \in \mathbb{R}^{N\times N}$, $\mathbf{n} = (n_1, n_2)$, $n_i = 1, ..., N$, for $m = (1, 2, 3)$, and $D_{\mathbf{F}}^{(0)}(0,\mathbf{n}) \triangleq \mathbf{F}$ [5]. 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 a scaling function $\psi^{(m)}$ with the coefficients $D_{\mathbf{F}}^{(0)}(j-1,\cdot)$, followed by decimation [1]. The wavelet leaders are then defined as the supremum of the normalized DWT coefficient magnitudes $d_{\mathbf{F}}^{(m)}(j,\mathbf{n}) \triangleq 2^{-j} D_{\mathbf{F}}^{(m)}(j,\mathbf{n})$ within the spatial neighborhood $3\lambda_{j,\mathbf{n}} = \bigcup_{i_1,i_2\in\{-1,0,1\}} \lambda_{j,(n_1+i_1,n_2+i_2)}$ over all finer scales:

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

where $\lambda_{j,\mathbf{n}}$ denotes a dyadic cube of side length 2^j at position $2^j\mathbf{n}$, i.e., $\lambda_{j,\mathbf{n}} = \{[2^j n_1, 2^j(n_1+1)), [2^j n_2, 2^j(n_2+1))\}$ [1].

Statistical model. Let $\ell_{j,\mathbf{n}} = \ln(L_{j,\mathbf{n}})$ be the log-leaders, the centered logarithms of the wavelet leader coefficients at scale j and at position $\mathbf{n}$, and $\boldsymbol{\ell} \triangleq [\boldsymbol{\ell}_{j_1}^T, ..., \boldsymbol{\ell}_{j_2}^T]^T$ the vector of concatenated log-leaders at scales $j_1, ..., j_2$, where $\boldsymbol{\ell}_j \in \mathbb{R}^{N_j^2}$, $N_j = N/2^j$, and $1 \leq j_1 \leq ... \leq j_2 \leq J$ (from fine to coarse resolution). According to [1], [6], the distribution of $\boldsymbol{\ell}_j$ can be approximated by a multivariate Gaussian distribution,

$$p(\boldsymbol{\ell}_j|\theta_1,\theta_2) \propto \left|\boldsymbol{\Sigma}_j(\boldsymbol{\theta},r)\right|^{-1/2} \exp\Big(-\frac{1}{2}\boldsymbol{\ell}_j^T \boldsymbol{\Sigma}_j(\boldsymbol{\theta},r)^{-1}\boldsymbol{\ell}_j\Big), \tag{2}$$

where $r \triangleq \|\Delta\mathbf{n}\|$, $|\mathbf{A}|$ is the matrix determinant, $\boldsymbol{\theta} \triangleq [\theta_1,\theta_2]^T$, $-c_2 = \theta_1$, θ_2 is related to the model adjustment constant c_2^0, and $\boldsymbol{\Sigma}_j(\boldsymbol{\theta},r) \approx \mathrm{Cov}(\boldsymbol{\ell}_{j,\mathbf{n}}, \boldsymbol{\ell}_{j,\mathbf{n}+\Delta\mathbf{n}})$ is the covariance defined by the radial symmetric function

$$\boldsymbol{\Sigma}_j(\boldsymbol{\theta},r) \triangleq \theta_1 g_1(j,r) + \theta_2 g_2(j,r), \tag{3}$$

where $g_1(j,r) = \max\{0, -\ln\big((r+1)/(r_j+1)\big)\}$ and $g_2(j,r) = \max\{0, -\ln(r+1)/\ln 4\}$ with $r_j = \lfloor N_j/4 \rfloor$ [6], [17]. Assuming independence between wavelet leaders at different scales yields

$$p(\boldsymbol{\ell}|\theta_1,\theta_2) = \prod_{j=j_1}^{j_2} p(\boldsymbol{\ell}_j|\theta_1,\theta_2). \tag{4}$$

Fourier domain likelihood. The direct use of the likelihood (4) in a numerical algorithm is difficult, but it can be efficiently approximated in the spectral domain by an equivalent Fourier-domain data augmented model introduced in [1]. This statistical model uses a Whittle approximation and data augmentation to facilitate the construction of conjugate priors for the parameters θ_1, θ_2 [1], [6]. Specifically, (4) is approximated as

$$p(\boldsymbol{\ell}|\theta_1,\theta_2) \approx p\big(\mathbf{x}, \boldsymbol{\mu}|\theta_1,\theta_2\big), \tag{5}$$

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

$$\begin{aligned} p\big(\mathbf{x}, \boldsymbol{\mu}|\theta_1,\theta_2\big) \propto\ & \theta_1^{-S} \exp\big(-\theta_1^{-1}(\mathbf{x}-\boldsymbol{\mu})^H G_{1,s}^{-1}(\mathbf{x}-\boldsymbol{\mu})\big) \\ & \times \theta_2^{-S} \exp\big(-\theta_2^{-1}\boldsymbol{\mu}^H G_{2,s}^{-1}\boldsymbol{\mu}\big). \end{aligned} \tag{6}$$

Here, $\boldsymbol{\mu}$ is a latent vector and $\mathbf{x} = [\mathbf{x}_1^T, ..., \mathbf{x}_S^T]^T$ is a subset of Fourier coefficients of $\boldsymbol{\ell}$ [1] defined on the regular grid $\mathcal{R}_j = [\![-N_j, N_j]\!]^2$, with $[\![a,b]\!]$ denoting the set of integers ranging from a to b, as

$$\mathbf{x}_s \triangleq \mathbf{x}(\boldsymbol{\omega}_s) = \frac{1}{N_j} \sum_{\mathbf{n}\in\mathcal{R}} \ell_{j,\mathbf{n}} \exp\left(-i\,\mathbf{n}^T\boldsymbol{\omega}_s\right), \tag{7}$$

where $s \in \big\{(j,\mathbf{m}) : j = [\![j_1, j_2]\!], \mathbf{m} \in \Omega_j^\dagger\big\}$, $\Omega_j^\dagger$ indexes low frequencies [1], and $\boldsymbol{\omega}_s = 2\pi\mathbf{m}/\sqrt{N_j}$ is the given Fourier frequency with $(j,\mathbf{m}) \in s$. The matrices $G_{i,s}$ in (6), for $i = 1, 2$, are real and positive and given by

$$G_{i,s} \triangleq \sum_{\mathbf{n}\in\mathcal{R}} g_i(j, \|\mathbf{n}\|) \exp\left(-i\,\mathbf{n}^T\boldsymbol{\omega}_\mathbf{s}\right). \tag{8}$$

III. Fourier Likelihood on Irregular Lattices

Although (6) enables fast estimation, its formulation relies on a regular spatial grid, making it unsuitable for heterogeneous image regions with irregular shapes and different MF properties. We develop a spectral likelihood capable of handling irregular lattices by considering debiasing constants [18]. Assume that the log-leaders $\boldsymbol{\ell}$ of an image $\mathbf{F}$ are a mixture of K classes $\mathcal{C}_1, ..., \mathcal{C}_K$ with a given MF parameter set per class: $\boldsymbol{\theta}_k \triangleq [\theta_1^k, \theta_2^k]$, for $k = 1, ..., K$. For log-leader segmentation purposes, consider the label vector $\mathbf{z} = [\mathbf{z}_{j_1}^T, ..., \mathbf{z}_{j_2}^T]^T$, where $\mathbf{z}_j = [z_{j,1}, ..., z_{j,N_j^2}]$ and $z_{j,n}$ denotes the label at position n and scale j that maps the observation $\ell_{j,n}$ to a given class $\mathcal{C}_k$.

To account for irregular domain shapes, we introduce an indicator function $\mathcal{J}_{j,\mathbf{n},k}$ defined as:

$$\mathcal{J}_{j,\mathbf{n},k} = \begin{cases} 1, & \mathbf{n} \in I_{j,k} = \{n \in N_j^2 \,|z_{j,n} = k\} \\ 0, & \text{otherwise,} \end{cases} \tag{9}$$

where $I_{j,k}$ is the subset of sites in $\mathbf{z}_j$ with label k at scale j. Using this notation, the likelihood (4) considering K MF classes is defined as:

$$p(\boldsymbol{\ell}|\boldsymbol{\theta}_1, ..., \boldsymbol{\theta}_K, \mathbf{z}) \triangleq \prod_{j=j_1}^{j_2} \prod_{k=1}^{K} p\big(\ell_{j,\{\mathbf{n}\in I_{j,k}\}}|\boldsymbol{\theta}_k\big), \tag{10}$$

where $p(\ell_{j,\{\mathbf{n}\in I_{j,k}\}}|\boldsymbol{\theta}_k)$ is the likelihood of a stationary zero mean Gaussian process with covariance (3) observed at locations $\mathbf{n} \in I_{j,k}$. Similar to (4), (10) poses challenges for the estimation of $\boldsymbol{\theta}_1, ..., \boldsymbol{\theta}_K$, and, deriving an approximation analogous to (6) for (10) is not as straightforward. Fortunately, by redefining (6) using the indicator function $\mathcal{J}$ for irregular lattices, we can completely remove estimation bias in spatial inference and tackle this problem [19].

Let $\mathbf{x} = [\mathbf{x}_1^T, ..., \mathbf{x}_K^T]^T$ be the debiased log-leader vector of Fourier coefficients with $\mathbf{x}_k = [\mathbf{x}_{1,k}, ..., \mathbf{x}_{S,k}]^T$, where

$$\mathbf{x}_{s,k} \triangleq \frac{1}{b_{j,k}} \sum_{\mathbf{n}\in\mathcal{R}} \mathcal{J}_{j,\mathbf{n},k}\big[\ell_{j,\mathbf{n}} - t_{j,k}\big] \exp\left(-i\,\mathbf{n}^T\boldsymbol{\omega}_s\right), \tag{11}$$

where $b_{j,k} = \sum_\mathbf{n} \mathcal{J}_{j,\mathbf{n},k}^{1/2}$ and $t_{j,k} = b_{j,k}\sum_\mathbf{n} \mathcal{J}_{j,\mathbf{n},k}\ell_{j,\mathbf{n}}$ are scaling constants for debiasing [19]. On the other hand, the covariance matrix in (6) needs to be adjusted with a weight

matrix $\mathbf{W}_{j,k}$ defined element-wise by

$$\mathbf{W}_{j,k}(n_1,n_2)=\frac{1}{b_{j,k}}\sum_{u,r\in\mathcal{R}_j}\mathcal{J}_{j,(u,r),k}\mathcal{J}_{j,(u+n_1,r+n_2),k}. \quad (12)$$

The augmented Whittle likelihood for k-th class then becomes

$$p\big(\mathbf{x}_k,\boldsymbol{\mu}_k|\theta_1^k,\theta_2^k,\mathbf{z}\big)\propto\prod_{s=1}^{S}(\theta_2^k)^{-S}\exp\Big(-\frac{1}{\theta_2^k}\boldsymbol{\mu}_k^H\mathbf{W}_{s,k}G_{2,s}^{-1}\boldsymbol{\mu}_k\Big)$$
$$\times(\theta_1^k)^{-S}\exp\Big(-\frac{1}{\theta_1^k}(\mathbf{x}_k-\boldsymbol{\mu}_k)^H\mathbf{W}_{s,k}G_{1,s}^{-1}(\mathbf{x}_k-\boldsymbol{\mu}_k)\Big), \quad (13)$$

where $\boldsymbol{\mu}_k=[\mu_{1,k}^T,...,\mu_{S,k}^T]^T$ and $\mathbf{W}_{s,k}\triangleq\mathbf{W}_{j,\mathbf{n},k}$. Similar to (5), we have:

$$p(\boldsymbol{\ell}|\boldsymbol{\Theta},\mathbf{z})\approx p\big(\mathbf{x},\boldsymbol{\mu}|\boldsymbol{\Theta},\mathbf{z}\big)\triangleq\prod_{k=1}^{K}p\big(\mathbf{x}_k,\boldsymbol{\mu}_k|\boldsymbol{\theta}_k,\mathbf{z}\big), \quad (14)$$

where $p\big(\mathbf{x},\boldsymbol{\mu}|\boldsymbol{\Theta},\mathbf{z}\big)$ is the spectral likelihood for all classes and $\boldsymbol{\Theta}=[\boldsymbol{\theta}_1,...,\boldsymbol{\theta}_K]$ is the vector of MF parameters.

IV. Bayesian Estimation and Gibbs Sampler

A. Prior and Posterior Distribution

Multiscale graph label prior $p(\mathbf{z}|\boldsymbol{\beta})$. The prior distribution of the labels $\mathbf{z}_j\in\{1,...,K\}^{N_j^2}$ is modeled with a modified version of the classical Potts-Markov random field (MRF)

$$p\big(\mathbf{z}_j|\mathbf{z}_{-j},\boldsymbol{\beta}_j\big)=\frac{1}{C(\boldsymbol{\beta}_j)}\exp\big[H(\mathbf{z}_j,\boldsymbol{\beta}_j)\big], \quad (15)$$

where $\mathbf{z}_{-j}$ indicates the labels on adjacent scales to j, i.e., $\mathbf{z}_{j-1}$, $\mathbf{z}_{j+1}$, H models variable relationships, and $\boldsymbol{\beta}_j=[\beta_{xy}^j,\beta_s]$ are the granularity parameters controlling the amount of spatial and inter-scale regularization, respectively. Moreover, $C(\boldsymbol{\beta}_j)\triangleq\sum_{\mathbf{z}_j}\exp[H(\mathbf{z}_j,\boldsymbol{\beta}_j)]$ is a partition function, which is in general intractable.

The prior (15) accounts for spatial and multiscale label correlations, which we propose to model with the function $H(\mathbf{z}_j,\boldsymbol{\beta}_j)$ defined as:

$$H(\mathbf{z}_j,\boldsymbol{\beta}_j)=\beta_{xy}^j\Phi(\mathbf{z}_j)+\beta_s\Psi(\mathbf{z}_j). \quad (16)$$

Here $\Phi(\mathbf{z}_j)$ models spatial dependencies and is defined as:

$$\Phi(\mathbf{z}_j)\triangleq\sum_{n=1}^{N_j^2}\sum_{m\in\mathcal{V}(n)}\delta\big(z_{j,n}-\mathbf{z}_{j,m}\big), \quad (17)$$

where $\mathcal{V}(n)$ contains the four spatial neighbors of n [20] δ is the Kronecker function, and

$$\Psi(\mathbf{z}_j)\triangleq\sum_{n=1}^{N_j^2}\sum_{\substack{p\in D(n)\\q\in U(n)}}\delta\big(z_{j,n}-\mathbf{z}_{j+1,p}\big)+\delta\big(z_{j,n}-\mathbf{z}_{j-1,q}\big) \quad (18)$$

models the inter-scale dependencies, where the functions $D(n)$ and $U(n)$ map the pixel n at a given scale to its corresponding parent or child coefficients, respectively. Fig. 1 (left) illustrates the proposed multiscale graph label prior.

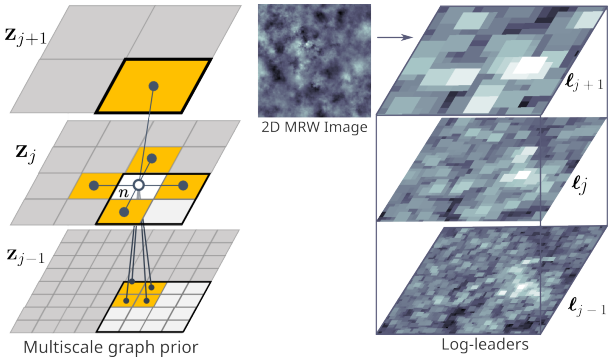

Fig. 1. Multiscale label graph structure (left) considered as a prior on the log-leader coefficients (left). The label at position n (white circle) uses its spatial and scale neighbors (gray dots) to update its value accordingly. This structure is proposed considering the hierarchical structure of the log-leaders, computed and plotted (right) from a 2D multifractal random walk (MRW).

Prior $p(\boldsymbol{\beta})$ for granularity parameters. We propose to also estimate the granularity parameters $\boldsymbol{\beta}$ using Gibbs sampling methods studied in [21], [22], circumventing the difficulties arising from the intractable constant $C(\boldsymbol{\beta}_j)$. Accordingly, a uniform prior distribution $\mathcal{U}_{(0,Q)}(\boldsymbol{\beta})$ is assumed for the granularity parameters on the interval $(0,Q)$, where $\boldsymbol{\beta}=[\beta_s,\beta_{xy}^{j_1},...,\beta_{xy}^{j_2}]$ gathers the whole set of parameters and Q is a given constant related to the number of classes (see [21] for details).

Prior $p(\boldsymbol{\Theta})$ for multifractal parameters. The natural conjugate priors for the multifractality parameters are inverse-gamma distributions, with parameters $\alpha_{i,k},\gamma_{i,k}$, for $i=1,2$, typically set to small values, ensuring that they closely resemble non-informative Jeffreys priors [1].

Posterior. The posterior distribution can be expressed as

$$p(\boldsymbol{\Theta},\boldsymbol{\mu},\mathbf{z},\boldsymbol{\beta}|\mathbf{x})\propto\ p\big(\mathbf{x},\boldsymbol{\mu}|\boldsymbol{\Theta},\mathbf{z}\big)\ p(\boldsymbol{\Theta})\ p(\mathbf{z}|\boldsymbol{\beta})\ p(\boldsymbol{\beta}), \quad (19)$$

where the spectral likelihood $p\big(\mathbf{x},\boldsymbol{\mu}|\boldsymbol{\Theta},\mathbf{z}\big)$ is given in (14). Unfortunately, estimators of the model parameters cannot be obtained in closed form. Instead, we propose to generate samples that are asymptotically distributed according to the target distribution using the Gibbs sampler [21]. Bayesian estimators of the unknown model parameters are then approximated using these samples.

B. Hybrid Gibbs Sampler and Bayesian estimators

The Gibbs sampler consists of generating samples from the conditional distributions $f(\boldsymbol{\Theta}|\mathbf{z},\boldsymbol{\beta},\mathbf{x},\boldsymbol{\mu})$, $p(\mathbf{z}|\mathbf{x},\boldsymbol{\mu},\boldsymbol{\Theta},\boldsymbol{\beta})$, and $f(\boldsymbol{\beta}|\boldsymbol{\Theta},\mathbf{x},\boldsymbol{\mu},\mathbf{z})=f(\boldsymbol{\beta}|\mathbf{z})$, where $f(\boldsymbol{\beta}|\mathbf{z})$ is challenging but can be approximated by a tractable sufficient statistic [21].

Sampling $\boldsymbol{z}$. The conditional distribution of a single label $z_{j,n}$, $j_1\leq j\leq j_2$, with all other labels $\mathbf{z}_{\backslash(j,n)}$ fixed is

$$z_{j,n}\,|\,\mathbf{z}_{\backslash(j,n)},\mathbf{x},\boldsymbol{\mu},\boldsymbol{\Theta},\boldsymbol{\beta}\propto p(z_{j,n}|\mathbf{z}_{\backslash(j,n)},\boldsymbol{\beta})p[\ell_{j,n}|z_{j,n}=k,\boldsymbol{\Theta}] \quad (20)$$

where $p[\ell_{j,n}|z_{j,n}=k,\boldsymbol{\Theta}]\sim\mathcal{N}(0,\sigma_j^2(\boldsymbol{\theta}_k))$ follows a zero-mean Gaussian distribution with variance $\sigma_j^2(\boldsymbol{\theta}_k)=\theta_2^k+\theta_1^k\log 2^j$, allowing (20) to be sampled efficiently using a checkerboard strategy exploiting conditional independence among labels.

Sampling $\boldsymbol{\Theta}, \boldsymbol{\mu}$. Given the labels, the MF parameters are sampled according to the conditional distributions

$$\begin{aligned}
\boldsymbol{\mu}_k|\mathbf{x}_k, \theta_1^k, \theta_2^k &\sim \mathcal{CN}\Big(\theta_1^k \tilde{G}_{1,k}\mathbf{x}_k, \big((\theta_1^k \tilde{G}_{1,k})^{-1} + (\theta_2^k \tilde{G}_{2,k})^{-1}\big)\Big), \\
\theta_1^k|\mathbf{x}_k, \boldsymbol{\mu}_k, \theta_2^k &\sim \mathcal{IG}\Big(\alpha_{1,k} + S, \gamma_{1,k} + \|\mathbf{x}_k - \boldsymbol{\mu}_k\|_{\tilde{G}_{1,k}}\Big), \\
\theta_2^k|\mathbf{x}_k, \boldsymbol{\mu}_k, \theta_1^k &\sim \mathcal{IG}\Big(\alpha_{2,k} + S, \gamma_{2,k} + \|\boldsymbol{\mu}_k\|_{\tilde{G}_{2,k}}\Big),
\end{aligned} \tag{21}$$

where $\|\mathbf{y}\|_{\mathbf{G}} \triangleq \mathbf{y}^H\mathbf{G}\mathbf{y}$, $\tilde{G}_{i,k} = \mathbf{W}_{\mathbf{s},k}G_{i,\mathbf{s}}^{-1}$, and $\mathcal{IG}$, $\mathcal{CN}$, respectively, denote inverse-gamma and complex Gaussian distributions.

Sampling $\boldsymbol{\beta}$. To generate samples for $\boldsymbol{\beta}$, we use the technique proposed in [21], [22], which performs gradient descent by approximating the intractable partition function $C(\boldsymbol{\beta}_j)$ with a given estimator. Steps 9-16 in Algo. 1 summarize the $\boldsymbol{\beta}$ sampling strategy, where the auxiliary vector $\mathbf{w}$ distributed according to the likelihood density $p(\mathbf{z}|\boldsymbol{\beta})$ is introduced for the gradient approximation (see [21], [22] for details). The statistics used here are the Potts potentials H, Φ and Ψ, defined in (16-18). The differences between the statistics evaluated in $\mathbf{z}$ and $\mathbf{w}$ are used to update the values of $\boldsymbol{\beta}$, which are constrained within $(0, Q)$ via $P_{[0,Q]}(x)$ (Steps 13, 16).

Bayesian estimators. Algorithm 1 summarizes the proposed hybrid Gibbs sampler, where "hybrid" arises from the heuristic approach for sampling the granularity parameters. The outputs of the Gibbs sampler are used to approximate the standard Bayesian estimators for MF the parameters (minimum mean square error (MMSE)) and labels (maximum a-posteriori (MAP)).

Algorithm 1: Hybrid Gibbs sampling procedure

Input :
$\mathbf{F} \in \mathbb{R}^{N\times N}$ (image); j_1, j_2 (initial,final scale); K (# regions); N_m (# iterations); N_b (burn-in period); N_v (# iterations for $\boldsymbol{\beta}$).

1 **Initialization:** $\mathbf{z}^{(0)}$, $\boldsymbol{\Theta}^{(0)}$, $\boldsymbol{\beta}^{(0)}$
2 Compute log-leaders $\boldsymbol{\ell}$ of $\mathbf{F}$
3 **for** $t = 1$ *to* N_m **do**
4 Compute $\mathbf{x}$ via (11)
5 Sample $\mathbf{z}^{(t)} \sim p\big[\mathbf{z}|\boldsymbol{\Theta}^{(t-1)}, \mathbf{z}^{(t-1)}, \boldsymbol{\beta}^{(t-1)}, \mathbf{x}\big]$ via (20)
6 Sample $\boldsymbol{\Theta}^{(t)} \sim f\big(\boldsymbol{\Theta}|\boldsymbol{\Theta}^{(t-1)}, \mathbf{z}^{(t)}, \boldsymbol{\beta}^{(t-1)}, \mathbf{x}\big)$ via (21)
7 **if** $t < N_b$ **then**
 ▷ Granularity parameter sampling
8 **for** $r = 1$ *to* N_v **do**
9 $\eta = 10\,(t + r - 1)^{-3/4}\big(\sum_{j=j_1}^{j_2} N_j\big)^{-1}$
10 **for** $j = j_1$ *to* j_2 **do**
11 Draw $\mathbf{w}_j \sim p[\mathbf{z}|\boldsymbol{\beta}^{(t)}]$
12 Set $\mathbf{q}_j = \Phi(\mathbf{z}_j) - \Phi(\mathbf{w}_j)$
13 $(\beta_{xy}^j)^{(r+1)} = P_{[0,Q]}\big((\beta_{xy}^j)_{(r)} + \eta\,\mathbf{q}_j\big)$
14 **end**
15 Set $\mathbf{w} \leftarrow [\mathbf{w}_{j_1}^T, ..., \mathbf{w}_{j_2}^T]^T$
16 $\beta_s^{(r+1)} = P_{[0,Q]}\big(\beta_s^{(r)} + \eta\big(H(\mathbf{z}, \boldsymbol{\beta}) - H(\mathbf{w}, \boldsymbol{\beta})\big)\big)$
17 **end**
18 Set $\boldsymbol{\beta}^{(t)} = \big[\beta_s^{(r+1)}, (\beta_{xy}^{j_1})^{(r+1)}, ..., (\beta_{xy}^{j_2})^{(r+1)}\big]$
19 **end**
20 **else**
21 Set $\boldsymbol{\beta}^{(t)} \leftarrow \boldsymbol{\beta}^{(t-1)}$ ▷ Fixed after burn-in period
22 **end**
23 **end**

V. Numerical Experiments

Monte Carlo simulation. To evaluate the proposed approach, we use 100 independent realizations of synthetic MF images (2D multifractal random walks (MRW) [1], [23]), containing an ellipse in the center with MF parameter $\theta_1^1 = -0.08$ and background with MF parameter $\theta_1^1 = -0.02$. A realization of a corresponding 2D MRW and its log-leaders is displayed in Fig. 1 (right) with the ground truth mask shown in Fig. 2. MRWs are non-Gaussian processes with stationary increments, designed to exhibit multifractal characteristics that resemble those of log-normal cascades of Mandelbrot [24].

The proposed method is compared against the k-means applied to patch-wise Bayesian MF estimation making use of a smoothing gamma MRF prior as in [1] (**K-GMRF**), and to Gabor filter features [25] (**K-GF**) (as an example of classical multiscale texture features). Moreover, a pretrained mask region-based convolutional neural network is used as a benchmark [26] (**Mask R-CNN**). The label performance is assessed using the dice score coefficient, DSC = 2TP/(2TP + FP + FN), and the classification error, Error = (FP + FN)/(TP + TN + FP + FN), based on true positive (TP), true negative (TN), false positive (FP), and false negative (FN) predictions. Estimations are made using a Daubechies mother wavelet with $N_\psi = 1$ vanishing moments, and with parameters set to $(\alpha_{i,k}, \gamma_{i,k}) = (10^{-3}, 10^{-3})$, $(j_1, j_2) = (1, 3)$, $K = 2$, $N_v = 2$, $N_m = 300$ iterations and burn-in period $N_b = 30$.

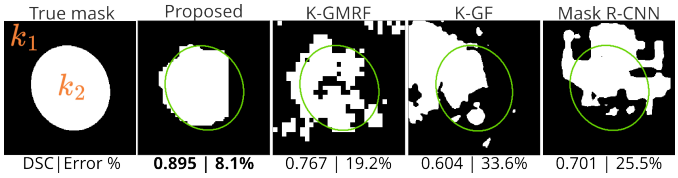

Fig. 2. Segmentation results using the compared methods of the 2D MRW from Fig. 1-right. The average DSC and error are shown for each result.

Segmentation performance. Fig. 2 shows the segmentation results for the single 2D MRW realization displayed in Fig. 1 (right), where the best segmentation performance is highlighted in bold. The proposed method yields an average DSC of 0.895 and the lowest error percentage, significantly surpassing the other approaches and providing better visual results. The average segmentation performance is given in Tab. I, confirming that the proposed method is interesting, achieving the highest segmentation scores and smallest classification error.

MF parameter estimation performance. Tab. II compares the average MF parameter estimation performance ([mean (standard deviation)] and root-mean square error (RMSE)) when labels are known (top) or estimated (bottom). The results

TABLE I
AVERAGE SEGMENTATION PERFORMANCE [MEAN (STD)].

Method	$\text{DSC}_{k=1}\uparrow$	$\text{DSC}_{k=2}\uparrow$	ERROR$\downarrow$ %
Proposed	**0.909** (0.034)	**0.752** (0.117)	**13.2** (4.8)
K-GMRF	0.843 (0.027)	0.567 (0.097)	22.8 (3.7)
K-GF	0.683 (0.084)	0.443 (0.126)	39.6 (8.4)
Mask R-CNN	0.813 (0.042)	0.300 (0.250)	28.5 (0.1)

TABLE II
AVERAGE MF PARAMETER ESTIMATION PERFORMANCE.

	k	True $\boldsymbol{\theta}_1^k$	MEAN (STD)	RMSE
labels known	1	-0.02	$-0.034\,(0.005)$	0.015
	2	-0.08	$-0.087\,(0.016)$	0.018
labels estimated	1	-0.02	$-0.033\,(0.006)$	0.014
	2	-0.08	$-0.113\,(0.025)$	0.041

show that in both cases that the estimation is very satisfactory, with overall slightly larger RMSE when labels are estimated, as expected. Moreover, results with known labels are similar to those obtained for single homogeneous textures [1], [6], confirming that the model proposed in Sec. III efficiently handles heterogeneous MF textures.

Overall, these results highlight that segmenting MF textures is a challenging task, even for trained models such as Mask R-CNN. Even so, the proposed method effectively identifies regions with different MF parameters and manages to estimate these parameters efficiently.

VI. CONCLUSIONS

This paper presented a Bayesian multifractal segmentation approach that models heterogeneous textures using wavelet leader multifractal analysis, a multiscale Potts prior and Fourier likelihood for irregular lattices. Our Bayesian framework efficiently computes the unknown model parameters using Gibbs sampling. Numerical experiments show the good performance of the proposed approach for image segmentation, surpassing local multifractality-based methods and a deep learning-based strategy. Future work includes extending the model to higher-dimensional models, as in [27], [28].

REFERENCES

[1] 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.

[2] R. Lopes and N. Betrouni, "Fractal and multifractal analysis: A review," *Medical Image Analysis*, vol. 13, no. 4, pp. 634–649, 2009.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] H. Wendt and P. Abry, "Multifractality tests using bootstrapped wavelet leaders," *IEEE Trans. Signal Proces.*, vol. 55, no. 10, pp. 4811–4820, 2007.

[9] 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.

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] 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.

[15] B. B. Chaudhuri and N. Sarkar, "Texture segmentation using fractal dimension," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 17, no. 1, pp. 72–77, 1995.

[16] Y. Xia, D. Feng, and R. Zhao, "Morphology-based multifractal estimation for texture segmentation," *IEEE Transactions on Image Processing*, vol. 15, no. 3, pp. 614–623, 2006.

[17] L. Leon, H. Wendt, J.-Y. Tourneret, and P. Abry, "A Bayesian framework for multivariate multifractal analysis," *IEEE Trans. Signal Proces.*, vol. 70, pp. 3663–3675, 2022.

[18] M. Fuentes, "Approximate likelihood for large irregularly spaced spatial data," *American Statistical Assoc.*, vol. 102, no. 477, pp. 321–331, 2007.

[19] 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.

[20] M. Pereyra, N. Dobigeon, H. Batatia, and J.-Y. Tourneret, "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.

[21] ——, "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.

[22] M. Pereyra, N. Whiteley, C. Andrieu, and J.-Y. Tourneret, "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.

[23] E. Bacry, J. Delour, and J.-F. Muzy, "Multifractal random walk," *Physical review E*, vol. 64, no. 2, p. 026103, 2001.

[24] B. Mandelbrot, "Intermittent turbulence in self-similar cascades: divergence of high moments and dimension of the carrier," *Multifractals and 1/ Noise: Wild Self-Affinity in Physics (1963–1976)*, pp. 317–357, 1999.

[25] A. K. Jain and F. Farrokhnia, "Unsupervised texture segmentation using Gabor filters," *Pattern Recognition*, vol. 24, no. 12, pp. 1167–1186, 1991.

[26] 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.

[27] 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.

[28] 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.