

Rigid-Motion-Invariant Classification of 3-D Textures

Saurabh Jain, Manos Papadakis, Sanat Upadhyay, and Robert Azencott

Abstract—This paper studies the problem of 3-D rigid-motion-invariant texture discrimination for discrete 3-D textures that are spatially homogeneous by modeling them as stationary Gaussian random fields. The latter property and our formulation of a 3-D rigid motion of a texture reduce the problem to the study of 3-D rotations of discrete textures. We formally develop the concept of 3-D texture rotations in the 3-D digital domain. We use this novel concept to define a “distance” between 3-D textures that remains invariant under all 3-D rigid motions of the texture. This concept of “distance” can be used for a monoscale or a multiscale 3-D rigid-motion-invariant testing of the statistical similarity of the 3-D textures. To compute the “distance” between any two rotations R_1 and R_2 of two given 3-D textures, we use the Kullback–Leibler divergence between 3-D Gaussian Markov random fields fitted to the rotated texture data. Then, the 3-D rigid-motion-invariant texture distance is the integral average, with respect to the Haar measure of the group $SO(3)$, of all of these divergences when rotations R_1 and R_2 vary throughout $SO(3)$. We also present an algorithm enabling the computation of the proposed 3-D rigid-motion-invariant texture distance as well as rules for 3-D rigid-motion-invariant texture discrimination/classification and experimental results demonstrating the capabilities of the proposed 3-D rigid-motion texture discrimination rules when applied in a multiscale setting, even on very general 3-D texture models.

Index Terms—Gaussian Markov random fields (GMRF), isotropic multiresolution analysis (IMRA), Kullback–Leibler (KL) divergence, rigid-motion invariance, volumetric textures, 3-D texture classification.

I. INTRODUCTION

OVER the course of the last two decades, a variety of deterministic or stochastic texture models and an even richer ensemble of texture discrimination methods have appeared in the literature, e.g., in [12], [15], [23], [34], and [45]. However, most of this work is exclusively devoted to 2-D textures. In this paper, we consider only stochastic 3-D textures, and we propose a novel method for 3-D rigid-motion-insensitive automatic texture discrimination. Although, throughout this paper, we exclusively refer to 3-D textures, the proposed approach to texture discrimination applies to 2-D stochastic

textures as well. In a nutshell, our contributions are: 1) the use of Isotropic Multiresolution Analysis (IMRA) to model novel 3-D rotation-covariant and 3-D translation-invariant texture signatures; and 2) the development of an efficiently computable distance and discriminant functions for 3-D rigid-motion-invariant texture discrimination/classification. For a detailed discussion of our contributions, we invite the reader to look at Section X. Section I is devoted to a brief survey of the literature and a heuristic presentation of our strategy to solve this problem. In Sections II through Sections VI, we develop the monoscale 3-D rigid-motion-invariant texture “distance.” Section VII discusses the multiscale 3-D rigid-motion-invariant texture discrimination. Section VIII describes how we construct 3-D texture realizations for our experiments, whereas in Section IX, we present our experiments and findings. In the same section, we show that the proposed methods can be used for the development of 3-D rigid-motion-invariant texture binary classification algorithms applicable to real-world 3-D images.

A. Stochastic Texture Models

Stochastic textures are often modeled by random fields. A random field $\{\mathbf{X}(s) : s \in S\}$ is a set of real-valued random variables defined on the same probability space. Such a random field is called Gaussian if any finite linear combination of $\{\mathbf{X}(s)\}$ has a Gaussian distribution. The first step in texture discrimination or classification is to define a texture signature, which is specified as a vector of computable texture features, belonging to a fixed finite or infinite dimensional feature vector space (FVS). The high dimensionality of FVS formally facilitates texture discrimination by their signatures at the expense of computational cost and higher sensitivity to noise. Past literature has been mostly focused on 2-D textures and has introduced a great variety of texture features ranging, for instance, from spatial frequencies based on Gabor filters [6], [18], [32], [40], [44] and wavelets [2], [10], [14], [20], [38], [46] to autoregressive models such as Markov random fields [5], [9], [11], [16], [26], [37]. For 2-D texture characterization by local texture information linked to multipixel statistics, spatial interaction stochastic models, such as Gibbs random fields, have been widely studied [17], [19], [22]. In particular, Gaussian Markov random fields (GMRFs) have been extensively applied for 2-D texture classification [11], [15], [27], [40], and, more recently, for 3-D textures in [41]. Existing literature on 3-D texture models and discrimination is still limited due to serious computational challenges native to 3-D image analysis (see [29], [30], [35], [42], and [48]).

Manuscript received January 15, 2011; revised August 06, 2011; accepted November 26, 2011. Date of publication January 27, 2012; date of current version April 18, 2012. This work was supported in part by the National Science Foundation under Grants NSF-DMS 0406748 and NSF-DMS 0915242. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Joseph P. Havlicek.

S. Jain is with the Center for Imaging Science, John Hopkins University, Baltimore, MD 21218 USA.

M. Papadakis, S. Upadhyay, and R. Azencott are with the Department of Mathematics of the University of Houston, Houston, TX 77204 USA (e-mail: mpapadak@math.uh.edu).

Digital Object Identifier 10.1109/TIP.2012.2185939

B. 3-D Rigid-Motion-Invariant Texture Discrimination

A natural assumption for textures is that their statistical properties of interest are invariant under translations. Therefore, the problem of 3-D rigid-motion-invariant texture discrimination reduces to that of 3-D rotationally invariant texture discrimination.

Background: Rotationally invariant discrimination between textures has been studied mostly in 2-D. For instance, in [28], two circularly symmetric autoregressive models are fitted to the 2-D texture data: one for the four nearest neighbors and one for the four diagonal nearest neighbors. Circularly symmetric models with higher order neighborhoods were used in [40], where the l^1 norm of the outputs of circularly symmetric 2-D Gabor filters was a key feature. A basic shortcoming of these models is that rotationally invariant discrimination between textures is achieved by handling nonisotropic textures as if they were isotropic. A mathematically rigorous treatment of isotropic textures can be found in [39]. However, generic textures often exhibit directional characteristics, which isotropic models cannot capture. The setback of these approaches is recognized in [15], where 2-D textures are modeled by continuous stationary GMRFs, and 2-D image texture rotations are generated via a continuous version of the discrete power spectral density defined by the digitized image data. The best rotation angle matching the given unknown texture to each texture from a predefined gallery of textures is estimated by maximizing a likelihood function derived from the continuous power spectral densities. Texture classification is then carried out by comparing model parameters. This approach worked well in 2-D, but a 3-D version would be computationally extremely expensive since the likelihood function in [15] is a product of terms evaluated at each node in a 3-D lattice. The approach of [21], based on local binary patterns, seems to be the first attempt to build a 3-D texture classification scheme robust to 3-D rotations.

Rationale and Outline of Our Approach: We work with arbitrary and generally nonisotropic 3-D textures, which we model as random fields that are defined on the continuous domain \mathbb{R}^3 . After image acquisition, the realizations of these textures are digital 3-D images, where we assume that image intensities at gray levels are given only for points in a discrete sampling lattice. Nonetheless, our approach easily extends to multicomponent images. Specifically, given an arbitrary spatially homogeneous 3-D texture \mathbf{X} with a zero mean that is defined on a fixed 3-D rectangular lattice $\mathbf{\Lambda} \subset \mathbb{R}^3$ similar to \mathbb{Z}^3 , we assume that \mathbf{X} is the restriction to $\mathbf{\Lambda}$ of a “continuous” 3-D texture that is defined on \mathbb{R}^3 , which we will also denote by \mathbf{X} for brevity. Obviously, the “continuous” parent texture is not affected by 3-D rigid motions. As we will explain later, we can focus only on 3-D rotations as 3-D shifts do not affect our texture model. However, the lattice of voxels $\mathbf{\Lambda} \subset \mathbb{R}^3$ is not invariant under 3-D rotations. Hence, for each 3-D rotation $\alpha \in SO(3)$, \mathbf{X}^α is defined by first applying the rotation α to the continuous “parent” texture \mathbf{X} and then by restricting this rotated continuous texture to $\mathbf{\Lambda}$. Moreover, the values of the continuous texture \mathbf{X} on the lattice points do not necessarily coincide with the values obtained during acquisition. To compatibly model the sampling of the realization of a continuous texture as the restriction of its values on a discrete lattice with the earlier natural definition for 3-D rotations of a discrete texture \mathbf{X} , we use the IMRA (see Section II) of its

“continuous” parent. The IMRA modeling of texture discretization facilitates at low computational cost a multiscale decomposition of the discrete texture \mathbf{X}^α at scales coarser than those of the original discrete texture \mathbf{X} . We fit to the rotated textures \mathbf{X}^α , at multiple scales j , zero-mean GMRFs $G_{\mathbf{X}}(\alpha, j)$ from the class introduced in [11] for 2-D and extended to 3-D in [41]. We choose this GMRF class because it facilitates a low computational cost parameter estimation, and the Kullback–Leibler (KL) distance between pairs of such GMRFs is explicitly computable (see Section IV-C). This low-dimensional model is obviously inadequate to describe complex and real-world textures. This is why, hereinafter, we use the term “fit the GMRF model to the texture” instead of “approximating the texture model.” However, in the experiment section, we demonstrate that the combined fitting of these simple models in a multiscale setting is capable of discriminating more complex and not necessarily GMRF textures. We next present a heuristic approach to incorporating rotational invariance in the proposed GMRF-fitting-based 3-D texture discrimination.

Given two discrete 3-D textures \mathbf{X} and \mathbf{Y} with zero intensity means, our preliminary choice for a texture “distance” at scale j could be the minimum taken over the $\tau \in SO(3)$ of the KL divergence between the GMRF models $G_{\mathbf{X}}(\mathbf{I}_3, j)$ and $G_{\mathbf{Y}}(\tau, j)$, where \mathbf{I}_3 is the 3×3 identity matrix. This model should be able, in principle, to give the best rotation matching angle between the two textures, but it does not lead to a rotationally invariant texture discriminant via fitting because the models $G_{\mathbf{X}}(\alpha, j)$ and $G_{\mathbf{Y}}(\tau\alpha, j)$ may not be the same as $G_{\mathbf{X}}(\mathbf{I}_3, j)$ and $G_{\mathbf{Y}}(\tau, j)$, respectively, as α runs throughout $SO(3)$. If a GMRF model is fitted to the texture \mathbf{X} , then a 3-D rotation of the same model can equally well be fitted to the same 3-D rotation of \mathbf{X} . In particular, the heart of the problem is not to classify textures as “identical” or “not identical” by finding their best matching angle but to develop a “distance” and, based on this “distance,” a decision function uniquely designating the value “identical” or “not identical” to all of their 3-D rigid motions. These considerations lead us to choose the orbit $\{G_{\mathbf{X}}(\alpha, j) : \alpha \in SO(3)\}$ of the fitted GMRF model of \mathbf{X} . Only those orbits are covariant to 3-D rotations and invariant to 3-D shifts. Thus, the problem that the first choice of texture distance appears to have is solved if we instead find a “distance” between the orbits $\{G_{\mathbf{X}}(\alpha, j) : \alpha \in SO(3)\}$ and $\{G_{\mathbf{Y}}(\tau\alpha, j) : \alpha \in SO(3)\}$. The average over $SO(3)$ of the KL divergences between the orbits for the fitted Gaussian fields $G_{\mathbf{X}}(\alpha, j)$ and $G_{\mathbf{Y}}(\alpha, j)$ is an obvious choice of such a “distance.” In fact, we define the 3-D rigid-motion-invariant distance between \mathbf{X} and \mathbf{Y} at the scale j as the minimum taken over $\tau \in SO(3)$ of the average of the KL divergences of the orbits of the GMRF models $G_{\mathbf{X}}(\alpha, j)$ and $G_{\mathbf{Y}}(\tau\alpha, j)$ as α traverses $SO(3)$ (see Section V-C).

II. IMRA FOR CONTINUOUS TEXTURES

Given an arbitrary 3-D texture, each 3-D rotation of this texture will be theoretically modeled below as a realization of a GMRF indexed by the Euclidean space \mathbb{R}^3 . For brevity, we will refer to such stochastic models as continuous textures. In practice, digital acquisition of 3-D textures delivers a discrete random field defined on a finite cube included in a discrete rectangular lattice, which can always be identified with $\mathbb{Z}^3 \subset \mathbb{R}^3$. We formalize the digitization process through a

multiscale approach implemented by MRA tools since textures can be intuitively viewed as a superposition of simpler textures “living” at different scales. Let D be the operator defined on $L^2(\mathbb{R}^3)$ by $Df(\mathbf{x}) = 2^{3/2}f(2\mathbf{x})$, and if $\mathbf{a} \in \mathbb{R}^3$, let $T_{\mathbf{a}}$ be given by $T_{\mathbf{a}}f(\mathbf{x}) = f(\mathbf{x} - \mathbf{a})$, where $\mathbf{x} \in \mathbb{R}^3$.

Definition 1: [43] An MRA of $L^2(\mathbb{R}^3)$ with respect to dyadic dilations is a sequence $\{V_j\}_{j \in \mathbb{Z}}$ of closed subspaces of $L^2(\mathbb{R}^3)$ satisfying the following conditions.

- 1) $\forall j \in \mathbb{Z}, V_j \subset V_{j+1}$.
- 2) $D^{-j}f$ belongs to V_0 if and only if f belongs to V_j .
- 3) $\cup_{j \in \mathbb{Z}} V_j$ is dense in $L^2(\mathbb{R}^3)$, and $\cap_{j \in \mathbb{Z}} V_j = \{0\}$.
- 4) There is a function ϕ such that V_0 is the closed linear span of $\{T_{\mathbf{n}}\phi : \mathbf{n} \in \mathbb{Z}^3\}$, which is a Bessel sequence.

This definition [8] extends the classical one [36] since we do not require $\{T_{\mathbf{n}}\phi : \mathbf{n} \in \mathbb{Z}^3\}$ to be an orthonormal or a Riesz basis of V_0 . Every V_j remains invariant under the action of the translation group $\{T_{2^{-j}\mathbf{n}} : \mathbf{n} \in \mathbb{Z}^3\}$. For high scales j , the space V_j is then practically translation invariant. However, even then, V_j is not invariant under the action of all 3-D rotations. Within this framework, the digitization at scale j of a continuous 3-D image f indexed by \mathbb{R}^3 is formalized by the linear mapping representing f by the sequence $\{\langle P_j f, D^j T_{\mathbf{n}}\phi \rangle\}_{\mathbf{n} \in \mathbb{Z}^3}$, where P_j is the orthogonal projection onto V_j . Each P_j commutes with translations $T_{2^{-j}\mathbf{n}}$ but not necessarily with arbitrary 3-D rotations. However, the resolution level required to produce satisfactory digital outputs should not change under image rotations. This simple pragmatic principle leads us to consider only MRAs for which each V_j is invariant under the action of the rotation group $SO(3)$. From here on, these MRAs will be called IMRAs for brevity. In this paper, we shall use IMRAs to generate multiscale representations for all 3-D rotations of continuous textures. Note that here and in the following, the Fourier transforms of any $f \in L^1(\mathbb{R}^3)$ are given by

$$\hat{f}(\xi) = \int_{\mathbb{R}^3} f(\mathbf{x})e^{-2\pi i(\mathbf{x} \cdot \xi)} d\mathbf{x}, \quad \xi \in \mathbb{R}^3.$$

Therefore, the Nyquist bandwidth is the cube $[-(1/2), (1/2)]^3$. The IMRAs generated by single refinable functions are characterized [43, Prop. 2.2] by the following necessary and sufficient condition: Let ϕ be a “refinable” function for which V_0 is the closed linear span of $\{T_{\mathbf{n}}\phi : \mathbf{n} \in \mathbb{Z}^3\}$. Then, $\{V_j\}_{j \in \mathbb{Z}}$ is an IMRA of $L^2(\mathbb{R}^3)$ if and only if ϕ is radial and $\hat{\phi}$ vanishes outside a ball $B(0, b)$ with $b < 1/2$.

III. AUTOCOVARANCE FUNCTION OF CONTINUOUS TEXTURES IN THE IMRA FRAMEWORK

Hereinafter, a continuous texture $\mathbf{X} = [\mathbf{X}_x], x \in \mathbb{R}^3$ is assumed to be generated by a stationary Gaussian random field. The autocovariance function ρ of \mathbf{X} is defined by $\rho(u) = \text{Cov}[\mathbf{X}_x, \mathbf{X}_{x+u}]$ for all u and x in \mathbb{R}^3 . We will always assume that:

- 1) The Fourier transform of ρ , which is denoted by $\hat{\rho}$, is C^∞ .
- 2) The support of $\hat{\rho}$ is contained in a ball centered at the origin.

Let $b_0 > 0$ be the radius of the ball in item 2 above, and let $0 < b < 1/2$ satisfy $2^{j-1}b < b_0 < 2^j b$ for some scale $j \in \mathbb{Z}^+$. We use an IMRA $\{V_j\}_j$ generated by a radial refinable function ϕ with a C^∞ Fourier transform $\hat{\phi}$ such that $\hat{\phi}(\xi) = 1$ for all

$|\xi| \leq 2^{-j}b_0$, and $\hat{\phi}(\xi) = 0$ for $|\xi| \geq b$. Then, V_0 contains all L^2 functions whose Fourier transform is supported in the ball of radius b ; therefore, $\rho \in V_j$. With no loss of generality, we may and will always assume that $j = 0$ so that $b/2 < b_0 < b$. Hence, $\hat{\phi} = 1$ on the support of $\hat{\rho}$, which implies that, for all $\mathbf{k} \in \mathbb{Z}^3$, $\langle \rho, T_{\mathbf{k}}\phi \rangle = \rho(\mathbf{k}) = \text{Cov}[\mathbf{X}_{\mathbf{k}}, \mathbf{X}_0]$. Note that the first of these two equalities essentially shows that the discrete sampling of a continuous texture is its restriction to a discrete lattice and that it requires the assumption $\hat{\phi} = 1$ on the support of $\hat{\rho}$. Therefore, the autocovariance of the discrete random field $\{\mathbf{X}(\mathbf{k}) : \mathbf{k} \in \mathbb{Z}^3\}$, which is the restriction of the continuous texture \mathbf{X} on \mathbb{Z}^3 , is simply the restriction of ρ to \mathbb{Z}^3 ; therefore, it is also denoted by ρ . In fact, the restriction of ρ to \mathbb{Z}^3 completely determines ρ on \mathbb{R}^3 via

$$\rho(\mathbf{s}) = \sum_{\mathbf{k} \in \mathbb{Z}^3} \rho(\mathbf{k})T_{\mathbf{k}}\phi(\mathbf{s}) \quad \forall \mathbf{s} \in \mathbb{R}^3. \quad (1)$$

The previous discussion establishes that the autocovariance function of a stationary stochastic 3-D texture \mathbf{X} indexed by \mathbb{R}^3 and that of its restriction to the infinite lattice \mathbb{Z}^3 are explicitly related via (1).

IV. CLASS OF GMRF MODELS

We emphasize that, in our approach, discrimination between two generic 3-D textures \mathbf{X} and \mathbf{Y} relies on the fitting of multiscale GMRF models after applying arbitrary rotations to \mathbf{X} and \mathbf{Y} . Fitting does not require \mathbf{X} or \mathbf{Y} to be Gaussian or stationary. Hence, parameter estimation of GMRFs is used only as a tool to extract local statistical characteristics of textures and is performed a fairly large number of times for each texture of interest. The complexity of these parameter estimations increases quickly when one increases the size of the basic neighborhoods defining these Gibbs random fields [31]. Thus, to minimize computing cost, we have deliberately restricted the class of GMRFs considered here to those with first-order neighborhoods. However, since such models essentially encode only correlations for nearby pixels, we first implement a multiscale decomposition of each texture \mathbf{X} by a fast IMRA decomposition algorithm (see Section VII-A), and we then fit GMRFs with first-order neighborhoods to each one of these monoscale texture outputs. This multiscale parametrization does encode medium- and long-distance correlations between voxels.

A. Stationary GMRF

We now describe the stationary GMRF models used in our approach (see [41]). We consider 3-D image data \mathbf{X} indexed by the finite lattice $\Lambda = \{\mathbf{k} = (i, j, k)\}$, where $1 \leq i \leq M$, $1 \leq j \leq N$, and $1 \leq k \leq P$. The grid points $\mathbf{k} = (i, j, k)$ are indexed by integers and are referred to as sites, nodes, or voxels. The lattice size is denoted by $|\Lambda| = MNP$. We assume that the image data $\mathbf{X}(\mathbf{k})$ belong to a finite set $\mathcal{A} \subset \mathbb{R}^+$ and that they are toroidal, i.e., they satisfy a periodic boundary condition. The spatial Markov structure of a stationary GMRF is defined by a translation-invariant neighborhood system $\eta = \{\eta_{\mathbf{k}}, \mathbf{k} \in \Lambda\}$ associating to each node \mathbf{k} the neighborhood $\eta_{\mathbf{k}} = (\mathbf{k} + W) \subset \Lambda$, where $W = W^+ \cup W^-$ is a symmetric neighborhood of the origin, i.e., $W^- = -W^+$. We

restrict our GMRF models to first-order interactions only, i.e., $W^+ = \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$.

The joint probability distribution of the random variables $\mathbf{X}(\mathbf{k})$, $\mathbf{k} \in \Lambda$ is a multivariate Gaussian distribution with covariance matrix $\Sigma = \sigma^2 \mathbf{B}^{-1}$ and mean vector $\boldsymbol{\mu}$. Stationarity forces the $|\Lambda| \times 1$ vector $\boldsymbol{\mu}$ to have constant coordinates $\boldsymbol{\mu}_{\mathbf{k}} = \mu$. The $(|\Lambda| \times |\Lambda|)$ matrix $\mathbf{B} = [b_{\mathbf{l}\mathbf{k}}]$ is positive and symmetric, and due to the assumed neighborhood structure W , we have

$$b_{\mathbf{l}\mathbf{k}} = \begin{cases} 1, & \text{if } \mathbf{l} = \mathbf{k}, \\ -\theta_{\mathbf{l}-\mathbf{k}} & \text{if } \mathbf{l} \in \eta_{\mathbf{k}}, \\ 0, & \text{else} \end{cases} \quad (2)$$

where $\theta_{\mathbf{r}} = \theta_{-\mathbf{r}}$ for all $\mathbf{r} \in W$. The matrix \mathbf{B} is parametrized by vector $\boldsymbol{\theta} = [\theta_{\mathbf{r}}]$, $\mathbf{r} \in W^+$. Hence, the GMRF models considered here are completely parametrized by μ , $\sigma > 0$, and the vector $\boldsymbol{\theta} = [\theta_{\mathbf{r}}]$, $\mathbf{r} \in W^+$ subject to the condition that $\mathbf{B} = \mathbf{B}(\boldsymbol{\theta})$ is a positive matrix. We assume that $\mu = 0$ for our textures.

The random ‘‘gray level’’ $\mathbf{X}(\mathbf{k})$ can be expressed by linear regression on the ‘‘gray levels’’ at neighboring nodes as follows [11], [27]:

$$x_{\mathbf{k}} = \sum_{\mathbf{r} \in W^+} \theta_{\mathbf{r}} (x_{\mathbf{k}-\mathbf{r}} + x_{\mathbf{k}+\mathbf{r}}) + e_{\mathbf{k}} \quad (3)$$

where $(e_1, \dots, e_{|\Lambda|})$ is the colored Gaussian noise with covariance $\sigma^2 \mathbf{B}(\boldsymbol{\theta})$.

Now, given a realization \mathbf{x} of the GMRF \mathbf{X} indexed by Λ , define for each $\mathbf{r} \in \Lambda$ the m -dimensional vector $\mathbf{y}_{\mathbf{r}} = [x_{\mathbf{l}} + x_{-\mathbf{l}}]$, $\mathbf{l} \in (\mathbf{r} + W^+)$, where m is the size of W^+ . Letting \mathbf{Y} to be the $(|\Lambda| \times m)$ matrix $\mathbf{Y} = [\mathbf{y}_{\mathbf{r}}]$, $\mathbf{r} \in \Lambda$, the least squares estimates [11] $\hat{\sigma}$ and $\hat{\boldsymbol{\theta}}$ of the parameters σ and $\boldsymbol{\theta}$ are given by the following statistics [41]:

$$\hat{\boldsymbol{\theta}}(\mathbf{x}) := (\mathbf{Y}^T \mathbf{Y})^{-1} \mathbf{Y}^T \mathbf{x} \quad (4)$$

$$\hat{\sigma}^2(\mathbf{x}) = \frac{1}{|\Lambda|} \left(\mathbf{x}^T \mathbf{x} - \hat{\boldsymbol{\theta}}^T \mathbf{Y}^T \mathbf{x} \right). \quad (5)$$

When W is a first-order neighborhood, as in our case, the positivity of the matrix $\mathbf{B}(\boldsymbol{\theta})$ is equivalent to the condition $|\boldsymbol{\theta}| < 0.5$, but in general, the condition for positivity is quite complicated [31]. Here, $|\boldsymbol{\theta}|$ denotes the ℓ^1 norm of vector $\boldsymbol{\theta}$. Hence, we restrict ourselves to this very simple neighborhood and perform the least squares estimate under the constraint $|\boldsymbol{\theta}| < 0.5$. Statistical properties of this estimate are analyzed in [27]. We choose this over maximum-likelihood (ML) schemes [11] because ML estimates require expensive numerical optimization techniques.

It is easy to check that the entries of $\mathbf{Y}^T \mathbf{Y}$ and $\mathbf{Y}^T \mathbf{x}$ can be determined from the empirical autocovariance function ρ_0 of the realization \mathbf{x} as follows:

$$\rho_0(\mathbf{k}) = \frac{1}{|\Lambda|} \sum_{\mathbf{r} \in \Lambda} x_{\mathbf{r}} x_{\mathbf{r}+\mathbf{k}}, \quad \text{for all } \mathbf{k} \in \Lambda. \quad (6)$$

The strictly stationary GMRF \mathbf{X} is assumed to have an autocovariance function ρ tending to 0 at infinity; hence, \mathbf{X} is ergodic. Thus, (see [4, Th. III.4.2 and 4.4]) for $|\Lambda|$ that is large enough, each $\rho(\mathbf{k})$ can be approximated with sufficient accuracy by the empirical autocovariance $\rho_0(\mathbf{k})$.

B. Extraction of Basic Texture Features

Let \mathbf{X} be an arbitrary digitized 3-D texture, approximately stationary, for which the acquisition process has delivered a realization \mathbf{x} indexed by a sufficiently large finite 3-D lattice Λ . We then systematically replace all $\mathbf{x}(\mathbf{k})$ by $\mathbf{x}(\mathbf{k}) - \hat{\mu}$, where $\hat{\mu}$ is the empirical mean of the realization \mathbf{x} . We then deliberately model \mathbf{x} exactly as if \mathbf{x} had been generated by a stationary GMRF associated to a neighborhood system of first order on the sublattice Λ , although this is generally not the case. This is what we refer to as ‘‘fitting the first-order GMRF model’’ on \mathbf{X} . We estimate the empirical variance σ^2 and the empirical covariances of \mathbf{x} , as well as the corresponding correlation parameter vector $\boldsymbol{\theta} \in \mathbb{R}^3$, using (4) and (5), under the necessary and sufficient condition $|\boldsymbol{\theta}| < 0.5$. In view of (2), the estimated variance $\hat{\sigma}^2$ and the estimated 3-D vector $\hat{\boldsymbol{\theta}}$ completely determine the estimated covariance matrix $\hat{\Sigma}(\mathbf{x})$ of the specific GMRF model that we have fitted to the 3-D texture realization \mathbf{x} . We essentially consider the matrix $\hat{\Sigma}(\mathbf{x})$ as a preliminary basic feature vector extracted from the realization \mathbf{x} . Further on, we will combine the preceding stationary GMRF modeling with 3-D rotations to generate high-dimensional signatures that are covariant under 3-D rigid motions.

C. KL Distance Between Stationary GMRFs

We will need to compare pairs of stationary GMRFs for discrimination purposes. A natural distance between stationary GMRFs can be derived from the KL divergence between stochastic processes. The advantages of the KL divergence over simpler Euclidean distances between parameter vectors, when dealing with stochastic models, are outlined in [6]. The KL divergence $\text{Div}(G_1 \| G_2)$ between two N -dimensional Gaussian probability distributions G_1 and G_2 , with respective mean vectors $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$ and invertible covariance matrices Σ_1 and Σ_2 , is given by

$$\begin{aligned} 2\text{Div}(G_1 \| G_2) &= \log_e \left(\frac{\det \Sigma_1}{\det \Sigma_2} \right) \\ &+ (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1)^T \Sigma_1^{-1} (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1) + \text{Trace} (\Sigma_1^{-1} \Sigma_2) - N. \end{aligned}$$

Symmetrizing the original KL divergence, which is not symmetric, and assuming zero means yield

$$\text{KLdist}(\Sigma_1, \Sigma_2) = (1/2) \text{Trace} (\Sigma_2^{-1} \Sigma_1 + \Sigma_1^{-1} \Sigma_2 - 2I_{N \times N}). \quad (7)$$

In our study, N is the size of the 3-D lattice Λ and is typically greater than 10^6 . For lattice sizes of this order, KLdist is not easily computable by means of the previous equation and with elementary matrix operations. However, since the GMRFs G_1 and G_2 satisfy a toroidal boundary condition and are strictly stationary, their covariance matrices commute with all periodic shifts in the following sense: Consider a tiling of \mathbb{Z}^3 by translates of Λ , i.e., there exists an infinite subset \mathcal{I} of \mathbb{Z}^3 such that \mathbb{Z}^3 is equal to the disjoint union of the tiles $\{\mathbf{i} + \Lambda : \mathbf{i} \in \mathcal{I}\}$. Then, for each $\mathbf{s} \in \mathbb{Z}^3$, $\mathbf{s} = \mathbf{i}(\mathbf{s}) + p(\mathbf{s})$ for unique choices of $\mathbf{i}(\mathbf{s}) \in \mathcal{I}$ and $p(\mathbf{s}) \in \Lambda$. Then, both covariance matrices are such that the (\mathbf{k}, \mathbf{l}) th entry of each Σ_i is equal to the $(p(\mathbf{s} + \mathbf{k}), p(\mathbf{s} + \mathbf{l}))$ th

entry of Σ_i . Matrices that satisfy this property are 3-D circulant convolution matrices. Standard arguments from operator theory or matrix analysis imply $\mathcal{F}(\Sigma_\kappa \mathbf{q}) = \mathbf{D}_\kappa(\mathcal{F}\mathbf{q})$, $\mathbf{q} \in \mathbb{R}^N$, $\kappa = 1, 2$, where \mathcal{F} is the 3-D DFT matrix and the \mathbf{D}_i values are $N \times N$ diagonal matrices; the entries of which are given by the restriction of the Fourier series of the autocovariance functions of the stationary GMRFs G_1 and G_2 , on a discrete lattice. Since the DFT is an isomorphism, we have

$$\text{KLdist}(\Sigma_1, \Sigma_2) = (1/2)\text{Trace}\left(\mathbf{D}_2^{-1}\mathbf{D}_1 + \mathbf{D}_1^{-1}\mathbf{D}_2 - 2\mathbf{I}_{N \times N}\right) \quad (8)$$

which is easily computable since inversion of diagonal matrices is a numerically inexpensive and stable operation. For the simple structure of the covariance matrix of our model, the entries of \mathbf{D}_i are samples of values of trigonometric polynomials directly obtained on points of the N^3 regular grid in the Fourier domain $[-(1/2), (1/2)]^3$. Thus, we do not require expensive fast Fourier transforms (FFTs) to calculate the distance.

V. ROTATIONALLY INSENSITIVE SIGNATURES FOR 3-D TEXTURES

A. 3-D Rotations of Textures

For all $\alpha \in SO(3)$ and $f \in L^2(\mathbb{R}^3)$, we define $\mathcal{R}_\alpha f(\mathbf{x}) = f(\alpha^T \mathbf{x})$. Similarly, the α rotation of a continuous 3-D texture \mathbf{X} indexed by \mathbb{R}^3 is the continuous texture where the random variable $(\mathbf{X}^\alpha)_\mathbf{s}$ corresponds to the random variable $\mathbf{X}_{\alpha^T \mathbf{s}}$. Then, the autocovariance ρ_α of the continuous texture \mathbf{X}^α is given by $\rho_\alpha = \mathcal{R}_\alpha \rho$. Our goal is to calculate the autocovariance of the discrete texture \mathbf{X}^α , which is the restriction of its continuous counterpart on the lattice \mathbb{Z}^3 . Equivalently, we want to estimate the sequence $\{\rho_\alpha(\mathbf{k})\}_{\mathbf{k} \in \mathbb{Z}^3}$ for all rotations α from the known input $\{\rho(\mathbf{k})\}_{\mathbf{k} \in \mathbb{Z}^3}$. This sequence completely determines the autocovariance ρ via (1). We assume that ρ belongs to V_0 of the IMRA generated by some radial function ϕ . This implies that ρ_α belongs to V_0 as well; hence, ρ_α is accurately represented by its values $\{\rho_\alpha(\mathbf{k})\}_{\mathbf{k} \in \mathbb{Z}^3}$ on the same lattice, regardless of the rotation α . Therefore, the \mathbf{k} th value of ρ_α is given by

$$\rho_\alpha(\mathbf{k}) = \langle \mathcal{R}_\alpha \rho, T_{\mathbf{k}} \phi \rangle = \langle \rho, T_{\alpha \mathbf{k}} \mathcal{R}_{\alpha^T} \phi \rangle = \langle \rho, T_{\alpha \mathbf{k}} \phi \rangle \quad (9)$$

for all $\mathbf{k} \in \mathbb{Z}^3$. Equation (9) proves a simple steerability rule, stating precisely how the sampling of a continuous texture by restriction of its values to the lattice \mathbb{Z}^3 is compatible with 3-D rotations. Moreover, if α_1 and α_2 are two rotations and if $(\rho_{\alpha_2})_{\alpha_1}$ is the autocovariance of the rotated texture \mathbf{X}^{α_2} by α_1 , then (9) implies

$$(\rho_{\alpha_2})_{\alpha_1}(\mathbf{k}) = \langle \rho_{\alpha_2}, T_{\alpha_1 \mathbf{k}} \phi \rangle = \langle \rho, T_{\alpha_2 \alpha_1 \mathbf{k}} \phi \rangle = \rho_{\alpha_2 \alpha_1}(\mathbf{k}) \quad (10)$$

which shows why (9) is a steerability rule. The computational implementation of this rule is done by approximating $\rho_\alpha(\mathbf{k}) = \langle \rho, T_{\alpha \mathbf{k}} \phi \rangle$ by $\langle \rho, T_{2^{-j_0} \mathbf{k}'} \phi \rangle$, where $\mathbf{k}' \in \mathbb{Z}^3$, by taking $j_0 > 0$ to be high enough so that points $2^{-j_0} \mathbf{k}'$ and $\alpha \mathbf{k}$ are sufficiently close. This computation is performed by iteratively applying j_0 steps of the reconstruction algorithm of the Fast IMRA transform [7], [25] on the data set $\{\rho_\alpha(\mathbf{k})\}_{\mathbf{k} \in \mathbb{Z}^3}$ in the following

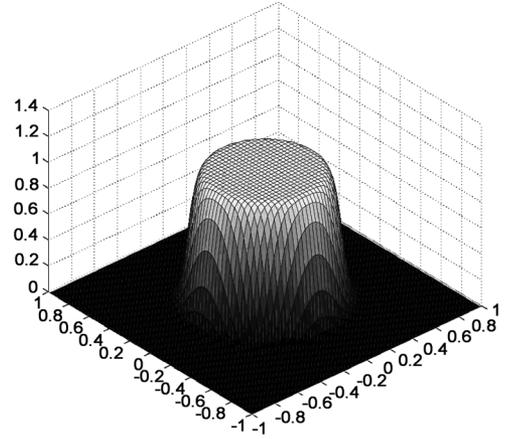


Fig. 1. 2-D plot of H_0 .

way: First, we define the transfer function of the low-pass synthesis filter H_0 as follows:

$$H_0(\xi) = \begin{cases} 1, & |\xi| \leq \frac{1-\beta}{2\tau} \\ \frac{1}{2} \left[1 + \cos\left(\frac{\pi\tau}{\beta} \left(|\xi| - \frac{1-\beta}{2\tau}\right)\right) \right], & \frac{1-\beta}{2\tau} < |\xi| \leq \frac{1+\beta}{2\tau} \\ 0, & \text{otherwise} \end{cases} \quad (11)$$

where $\xi \in [-(1/2), (1/2)]^3$, β and $\tau > 0$, and H_0 is \mathbb{Z}^3 periodic. The parameter $\beta > 0$ determines the width β/τ of the transition band, whereas τ determines the cutoff frequency of the filter (see Fig. 1). Typically, we set $\tau = 100/84$ and $\beta = 1/7$. At each reconstruction step, the low-pass component at the current resolution level $j < j_0$ is first upsampled by a factor of 2. The upsampled low-pass component is convolved with the synthesis low-pass filter h_0 . It is best to implement this convolution as fast convolution in the frequency domain via multiplication with H_0 . The added benefit is speed and best use of the isotropy of H_0 , which manifests better in the frequency domain. At the first iteration of the reconstruction algorithm, the low-pass component is set to a given input image. For instance, for the implementation of (10), we use $\{\rho_\alpha(\mathbf{k})\}_{\mathbf{k} \in \mathbb{Z}^3}$. Then, the current high-pass component, which is forced to be zero, is added to generate the image at resolution level $j + 1$. In practice, all this computation is restricted on a finite sublattice of \mathbb{Z}^3 . In Section IX, we use $j_0 = 1, 2$. We also observe that this method of approximating $\rho_\alpha(\mathbf{k})$ is numerically more accurate than trilinear interpolation between nearest neighbors of $\alpha \mathbf{k}$.

B. Monoscale Rotationally Covariant Texture Signatures

Let \mathbf{x} be a realization of the discrete stationary zero-mean 3-D texture \mathbf{X} indexed by a sufficiently large finite 3-D lattice $\Lambda \subset \mathbb{Z}^3$. As before, ρ_α ($\alpha \in SO(3)$) is the autocovariance function of \mathbf{X}^α that is calculated as described in the Section V-A, where ρ is estimated by the empirical autocovariance ρ_0 by means of (6). Under the assumption that the given discrete texture \mathbf{X} at the scale that corresponds to the grid Λ has zero mean and is wide-sense stationary, we calculate its empirical autocovariance

matrix $\hat{\Sigma}$, following Section IV-B. In light of (10), we define the texture signature of \mathbf{X} by

$$\Gamma_{\mathbf{X}}(\boldsymbol{\alpha}) = \Sigma_{\boldsymbol{\alpha}} \quad \boldsymbol{\alpha} \in SO(3) \quad (12)$$

where $\Sigma_{\boldsymbol{\alpha}}$ is the autocovariance of $\mathbf{X}^{\boldsymbol{\alpha}}$. The strict stationarity of \mathbf{X} and of all $\mathbf{X}^{\boldsymbol{\alpha}}$ imply that the (\mathbf{k}, \mathbf{l}) entry of $\Sigma_{\boldsymbol{\alpha}}$ is equal to $\rho_{\boldsymbol{\alpha}}(\mathbf{k} - \mathbf{l})$. By combining this choice of texture signature with (10), we obtain $\Gamma_{\mathbf{X}^{\boldsymbol{\alpha}_2}}(\boldsymbol{\alpha}_1) = \Gamma_{\mathbf{X}}(\boldsymbol{\alpha}_2 \boldsymbol{\alpha}_1)$, for all $\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2 \in SO(3)$; i.e., $\Gamma_{\mathbf{X}}$ is a rotationally covariant signature for the 3-D texture \mathbf{X} . This is a key property in the construction of rotationally invariant distances presented in Section V-C. However, in practice, Σ is not available. Instead, we obtain from \mathbf{x} the empirical autocovariance ρ_0 , which is an estimate of ρ . Then, for each $\boldsymbol{\alpha} \in SO(3)$, we derive the values of $\{(\rho_0)_{\boldsymbol{\alpha}}(\mathbf{k})\}_{\mathbf{k} \in \Lambda}$ from those of ρ_0 according to the previous subsection. We define

$$\Gamma_{\mathbf{x}}(\boldsymbol{\alpha}) = (\hat{\Sigma})_{\boldsymbol{\alpha}} \quad \boldsymbol{\alpha} \in SO(3) \quad (13)$$

where $(\hat{\Sigma})_{\boldsymbol{\alpha}}$ is the autocovariance matrix corresponding to $\{(\rho_0)_{\boldsymbol{\alpha}}(\mathbf{k})\}_{\mathbf{k} \in \Lambda}$ to be the monoscale 3-D texture signature of the observed texture \mathbf{X} at the scale corresponding to the density of the lattice Λ . This is the orbit $\{G_{\mathbf{X}}(\boldsymbol{\alpha}, j) : \boldsymbol{\alpha} \in SO(3)\}$ of the GMRF models fitted to the texture \mathbf{X} and to all of its rotations at scale $j = 0$ that we promised in the introduction. Now, let $\boldsymbol{\alpha}_1 \in SO(3)$. Let \mathbf{x}, \mathbf{y} be the realizations of \mathbf{X} and $\mathbf{X}^{\boldsymbol{\alpha}_1}$ observed on Λ , and let $\Gamma_{\mathbf{x}}$ and $\Gamma_{\mathbf{y}}$ be their signatures, respectively. If $\{\rho(\mathbf{k})\}_{\mathbf{k} \in \mathbb{Z}^3}$ is the 3-D autocovariance of \mathbf{X} , then the 3-D autocovariance of $\mathbf{X}^{\boldsymbol{\alpha}_1}$ is $\{\rho_{\boldsymbol{\alpha}_1}(\mathbf{k})\}_{\mathbf{k} \in \mathbb{Z}^3}$. Ideally, one should expect then that the empirical autocovariance of $\mathbf{X}^{\boldsymbol{\alpha}_1}$ estimated from \mathbf{y} should be equal to $\{(\rho_0)_{\boldsymbol{\alpha}_1}(\mathbf{k})\}_{\mathbf{k} \in \Lambda}$. Since empirical autocovariances are only approximations of true autocovariances, this equality can be only true within a certain margin of error, but this is not the only reason why the rotational covariance of $\Gamma_{\mathbf{x}}$ only holds approximately. The reason for which $\Gamma_{\mathbf{y}}(\boldsymbol{\alpha}) \approx \Gamma_{\mathbf{x}}(\boldsymbol{\alpha}_1 \boldsymbol{\alpha})$, for all $\boldsymbol{\alpha} \in SO(3)$, is that the computation of $(\rho_0)_{\boldsymbol{\alpha}_1}(\mathbf{k})$ is implemented via an approximation. If there had been no approximation in the computation of these coefficients, $\Gamma_{\mathbf{y}}(\boldsymbol{\alpha}) = \Gamma_{\mathbf{x}}(\boldsymbol{\alpha}_1 \boldsymbol{\alpha})$, for all $\boldsymbol{\alpha} \in SO(3)$, would be true. Therefore, the faithfulness of the rotational covariance of the texture signatures that we compute lies entirely on the global approximation of the sequence $\{(\rho_0)_{\boldsymbol{\alpha}_1}(\mathbf{k})\}_{\mathbf{k} \in \Lambda}$ for each desired rotation $\boldsymbol{\alpha}_1$. We now introduce rotationally invariant monoscale 3-D texture distances.

C. Rotationally Invariant Distance Between 3-D Texture Signatures

Recall that the KL distance between two stationary Gaussian random fields with zero means and autocovariance matrices Σ_1 and Σ_2 is a symmetric function $\text{KLdist}(\Sigma_1, \Sigma_2)$ given by (7). The Markov property of GMRFs is critical for a low computational cost estimation of the inverses of Σ_1 and Σ_2 . Consider a 3-D texture \mathbf{X} , and let $\mathbf{X}^{\boldsymbol{\beta}}$ be the rotation of \mathbf{X} induced by $\boldsymbol{\beta} \in SO(3)$. The corresponding signatures $\Gamma_{\mathbf{X}}$ and $\Gamma_{\mathbf{X}^{\boldsymbol{\beta}}}$ defined by (12) are both indexed by $SO(3)$ and differ only by a multiplicative “shift” acting on their index space, namely multiplication

by $\boldsymbol{\beta}$. We want these textures to be classified as identical. Since direct comparison of signatures indexed by $SO(3)$ is quite impractical, we introduce an “average divergence” between pairs of texture signatures: For any pair \mathbf{X}, \mathbf{Y} of 3-D textures with autocovariance functions satisfying the mild properties stated in Section III, we will then define a distance $\text{dist}(\mathbf{X}, \mathbf{Y})$ by

$$\text{dist}(\mathbf{X}, \mathbf{Y}) := \min_{\boldsymbol{\tau} \in SO(3)} \int_{SO(3)} \text{KLdist}(\Gamma_{\mathbf{X}}(\boldsymbol{\alpha}), \Gamma_{\mathbf{Y}}(\boldsymbol{\tau} \boldsymbol{\alpha})) d\boldsymbol{\alpha}. \quad (14)$$

Since the integral in (14) is continuous in $\boldsymbol{\tau}$ over the compact group $SO(3)$, it must achieve its minimum for at least one rotation $\boldsymbol{\tau} \in SO(3)$, which is then one of the rotations that best aligns the signatures of \mathbf{X} and $\mathbf{Y}^{\boldsymbol{\tau}}$.

The left invariance of the Haar measure [13] and the preceding definition readily imply the main rotation invariance properties as follows:

$$\text{dist}(\mathbf{X}, \mathbf{Y}) = \text{dist}(\mathbf{X}^{\boldsymbol{\alpha}}, \mathbf{Y}^{\boldsymbol{\beta}}) \quad \forall \boldsymbol{\alpha}, \boldsymbol{\beta} \in SO(3) \quad (15)$$

$$\text{dist}(\mathbf{X}, \mathbf{X}^{\boldsymbol{\alpha}}) = 0 \quad \forall \boldsymbol{\alpha} \in SO(3). \quad (16)$$

As stated earlier, in practice, the original texture signatures $\Gamma_{\mathbf{X}}$ and $\Gamma_{\mathbf{Y}}$ are not available and must be replaced by 3-D texture signatures that are derived from realizations \mathbf{x} and \mathbf{y} of these 3-D textures given on a finite lattice Λ . This leads us to compute a distance $\text{Rdist}(\mathbf{x}, \mathbf{y})$ approximating $\text{dist}(\mathbf{X}, \mathbf{Y})$ and naturally defined by

$$\text{Rdist}(\mathbf{x}, \mathbf{y}) := \min_{\boldsymbol{\tau} \in SO(3)} \int_{SO(3)} \text{KLdist}(\Gamma_{\mathbf{x}}(\boldsymbol{\alpha}), \Gamma_{\mathbf{y}}(\boldsymbol{\tau} \boldsymbol{\alpha})) d\boldsymbol{\alpha}. \quad (17)$$

Remark 1: When the diameter of the lattice Λ tends to infinity, then $\text{Rdist}(\mathbf{x}, \mathbf{y})$ converges to $\text{dist}(\mathbf{X}, \mathbf{Y})$. The rotation invariance of $\text{Rdist}(\mathbf{x}, \mathbf{y})$ also holds approximately because the texture signatures are only approximately covariant as we have explained in the previous section. On the other hand, $\text{Rdist}(\mathbf{x}, \mathbf{y})$ is not influenced by the translates of \mathbf{x} and \mathbf{y} by any vector \mathbf{t} as long as \mathbf{t} belongs to Λ because texture signatures are computed by approximate autocovariance functions, which themselves are invariant to translates that leave the sampling lattice Λ invariant.

Hence, in all practical tests and applications that follow, we consider $\text{Rdist}(\mathbf{x}, \mathbf{y})$ to be the promised 3-D rigid-motion-invariant texture “distance.”

VI. NUMERICAL COMPUTATION OF ROTATIONALLY INVARIANT DISTANCE BETWEEN 3-D TEXTURES

A. Parametrization of $SO(3)$ by Euler Angles

Each 3-D rotation $\boldsymbol{\alpha} \in SO(3)$ is parametrized by its three Euler angles (α, β, γ) , using the ZYZ convention, and can be written as the product of three rotations as follows:

$$\boldsymbol{\alpha} = R_Z(\gamma)R_Y(\beta)R_Z(\alpha) \quad (18)$$

where $R_Z(\gamma)$, $R_Y(\beta)$, and $R_Z(\alpha)$ are rotations by angles γ, β, α around the Z -axis, the Y -axis, and the Z -axis, respectively, with $0 \leq \alpha, \gamma < 2\pi$, and $0 \leq \beta \leq \pi$.

The normalized left-invariant Haar measure $SO(3)$ is then given by

$$d\boldsymbol{\alpha} = \frac{1}{8\pi^2} \sin(\beta) d\alpha d\beta d\gamma \quad (19)$$

where $d\alpha$, $d\beta$, and $d\gamma$ are the usual Lebesgue measure.

Due to the symmetries of covariance matrices in our texture models, the integral in (17) may be restricted to Euler angles as follows:

$$(\alpha, \beta, \gamma) \in EB = [0, \pi] \times \left[0, \frac{\pi}{2}\right] \times [0, \pi]. \quad (20)$$

Indeed, for $\alpha = a + \pi$, with $0 < a < \pi$, $0 \leq \beta \leq (\pi/2)$, $0 \leq \gamma \leq \pi$, we have

$$\boldsymbol{\alpha} = \boldsymbol{\alpha}_1 R_Z(\pi)$$

where $\boldsymbol{\alpha}_1 = R_Z(\gamma)R_Y(\beta)R_Z(a)$. Since our covariance matrices are invariant under the rotation $R_Z(\pi)$, we can take $\alpha \in [0, \pi]$. Similarly, γ can be restricted to $[0, \pi]$ as well.

To justify why it is enough to take β in $[0, (\pi/2)]$, first, let $\boldsymbol{\alpha} = R_Z(\gamma)R_Y(a + \pi/2)R_Z(\alpha)$ with $0 < a < (\pi/2)$. Now, $R_Y(a + (\pi/2))$ can be replaced by $R_Y(a + (3\pi/2))$ because the latter rotation is the composition of rotations $R_Y(\pi)$ and $R_Y(a + (\pi/2))$, and the covariance matrices we consider are invariant under $R_Y(\pi)$. To complete the proof of this claim note

$$R_Y\left(a + \frac{3\pi}{2}\right) = R_Z(\pi)R_Y\left(\frac{\pi}{2} - a\right)R_Z(\pi).$$

B. Discretization of $SO(3)$

For the practical computation of the rotationally invariant distance $Rdist$ defined in (17), we discretize the integral over $SO(3)$ in the right-hand side of (17).

In view of Section VI-A, we select a finite 3-D rectangular grid of points $(\alpha, \beta, \gamma) \in EB$ by defining a partition of EB into rectangles having the same Haar measure. In view of (19), this is achieved by defining

$$\begin{aligned} \alpha_i &= \left\{ \frac{i\pi}{N} \right\} \quad \forall i = 0, 1, \dots, N-1 \\ \beta_j &= \left\{ \arccos\left(1 - \frac{j+0.5}{N}\right) \right\} \quad \forall j = 0, 1, \dots, N-1 \\ \gamma_k &= \left\{ \frac{k\pi}{N} \right\} \quad \forall k = 0, 1, \dots, N-1 \\ \boldsymbol{\alpha}_{(i,j,k)} &= (\alpha_i, \beta_j, \gamma_k). \end{aligned}$$

The discrete values of β start with 0.5 to avoid the unwieldy gimbal lock point $\beta_0 = 0$ (see [1]). Using the discrete set of rotations that were just defined, a computationally implementable version of the distance $Rdist$ defined in (17) is given by

$$\min_{\boldsymbol{\tau} \in SO(3)} \frac{1}{N^3} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} \text{KLdist}\left(\Gamma_{\mathbf{x}}(\boldsymbol{\alpha}_{(i,j,k)}), \Gamma_{\mathbf{y}}(\boldsymbol{\tau}\boldsymbol{\alpha}_{(i,j,k)})\right). \quad (21)$$

By taking into account the symmetries of the covariance matrices, the numerical minimization in (21) can be restricted to $\boldsymbol{\tau} \in EB$. Thus, computing the approximate $Rdist$ defined by

(21) involves solving the following minimization problem in three variables:

$$\text{Find } \min_{\boldsymbol{\tau} \in EB} F(\boldsymbol{\tau}) \quad (22)$$

where $F(\boldsymbol{\tau})$ is

$$\frac{1}{N^3} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} \text{KLdist}\left(\Gamma_{\mathbf{x}}(\boldsymbol{\alpha}_{(i,j,k)}), \Gamma_{\mathbf{y}}(\boldsymbol{\tau}\boldsymbol{\alpha}_{(i,j,k)})\right).$$

To keep the computations tractable for intensive numerical testing on a set of roughly 6400 pairs of standard-size 3-D images, we choose $N = 5$, which results in a total of 125 distinct rotations; thus, the texture signature involves $125 \times 4 = 500$ parameters. For smaller test sets, one could of course use $N > 5$. Each evaluation of F requires the rotation of one of the two textures: the calculation of 125 sets of GMRF parameters for the rotated texture and the calculation of KL distance for 125 pairs of parameters.

The rotation of a texture is achieved by rotating the covariance matrix, as explained earlier in Section V-A, which is a computationally much cheaper task than rotating the entire 3-D texture. For efficiency, each KL distance is calculated in the Fourier domain. Thus, each evaluation of the function $F(\boldsymbol{\tau})$ is achieved in about 1 s, which is still quite costly for one function evaluation. For instance, finding the minimum of F by exhaustive search through 1000 rotations (corresponding to $N = 10$) would result in a computation time of 15 min. An optimization approach involving derivatives would also require a high number of evaluations of F and might get trapped in local minima.

To circumvent these problems, we minimize $F(\boldsymbol{\tau})$ through the Matlab pattern-search algorithm [3], which yields a good minimum of F in less than 100 function evaluations. On a 2.8-GHz processor, the calculation of $Rdist$ between two given 3-D textures of size $120 \times 120 \times 120$ takes 2–3 min.

VII. APPLICATION OF 3-D RIGID-MOTION-INVARIANT TEXTURE DISCRIMINATION

A. Multiscale Rotationally Invariant Distances of 3-D Textures

The computable distance $Rdist$ quantifies the dissimilarity between two 3-D textures given on a finite lattice Λ . One of the main ingredients of our approach is the assumption that there is an IMRA $\{V_j\}_j$ generated by a radial function ϕ such that the autocovariance ρ of the texture \mathbf{X} belongs to a subspace V_j . The scale j determines the density of the lattice Λ . By convention, we can always set $j = 0$. However, we do not need to confine the $Rdist$ -based comparison of two textures to a single-scale $j = 0$ since the $Rdist$ -based comparison of two textures at other scales probes into correlations of voxels that are further apart and therefore enhances texture discrimination.

Using this multiscale comparison of textures, we first gain reduced computational cost at each scale because the first-order GMRF parametrization model we rely on is small. This would definitely not be the case if we had decided to use “order- k ” GMRF models with $k \gg 1$ (see [31]). Second, GMRF models with $k > 1$ are determined by a parameter vector $\boldsymbol{\theta}$ for which the constraint $|\boldsymbol{\theta}| < 1/2$ is only sufficient and severely restricts the parameter space. To implement the multiscale comparison

of two textures, we use the IMRA decomposition algorithm [7], [39] in the following way: The input image is convolved with analysis low-pass filter h_0 first, and then, it is downsampled by a factor of 2. This procedure produces the low-pass output, which can be used as input for an iteration of the low-pass decomposition. The transfer function of the analysis low-pass filter is given by (11) with $\tau = 100/59$ and $\beta = 1/6$. The convolution by h_0 can again be performed as fast convolution in the frequency domain. If we wish to maintain translation invariance, then we apply the decomposition algorithm, and we skip the downsampling by 2. In this case, for a single application of the decomposition algorithm, the first-order neighbors are the third-order neighbors, i.e., two voxels apart on each of the x , y , and z -directions.

To implement the multiscale Rdist-based comparison of two textures, we begin with two realizations $\mathbf{x}^{(0)}$ and $\mathbf{y}^{(0)}$ of the arbitrary 3-D textures \mathbf{X} and \mathbf{Y} given above. Both $\mathbf{x}^{(0)}$ and $\mathbf{y}^{(0)}$ are given on a finite lattice $\Lambda^{(0)}$. We apply the IMRA decomposition algorithm to $\mathbf{x}^{(0)}$ and $\mathbf{y}^{(0)}$ to generate two corresponding multiscale realizations $\mathbf{x}^{(j)}$ and $\mathbf{y}^{(j)}$ defined for each of the coarser scales $j = -1, -2, -3, \dots, -J$, on successively coarser lattices $\Lambda^{(j)}$. At each coarser scale $j-1$, the lattice $\Lambda^{(j)}$ is reduced by a factor of two per dimension. Hence, the number $J+1$ of scales that we can consider in 3-D is fairly small in practice and depends essentially on the initial resolution of our data, which is the cardinality of the initial lattice $\Lambda^{(0)}$.

Then, for $j = 0, -1, -2, -3, \dots, -J$, we define $\text{Rdist}_j(\mathbf{x}, \mathbf{y}) = \text{Rdist}(\mathbf{x}^{(j)}, \mathbf{y}^{(j)})$. The vector of these $J+1$ rotationally invariant distances $\text{Rdist}_j(\mathbf{x}, \mathbf{y})$, where $j = 0, -1, \dots, -J$, will be our multiscale 3-D texture discrimination tool.

B. Self-Distance for 3-D Textures

Theoretically, the self-distance $\text{Rdist}_j(\mathbf{x}, \mathbf{x}^\alpha)$ of the 3-D texture \mathbf{x} from every rotation of itself at scale j must be zero. However, when we compute the distance of a 3-D texture from a rotation of itself, we only achieve an approximate minimization in formula (22). Furthermore, note that we implement texture rotations via a rotation of their empirical autocovariance function, which adds another level of approximation to $\text{Rdist}_j(\mathbf{x}, \mathbf{x}^\alpha)$. Hence, numerically computed “self-distances” are small but nonzero. Therefore, discrimination between 3-D textures must take these nonzero “self-distances” into account.

To estimate the self-distance of a 3-D texture, we generate rotations of $\mathbf{X}^{(j)}$ by 20 randomly selected $\alpha_k \in SO(3)$ ($k = 1, 2, \dots, 20$) and define $\text{diam}_\mathbf{x}(j)$ to be the 80th percentile of the 20 distances $\text{Rdist}(\mathbf{x}^{(j)}, (\mathbf{x}^{(j)})^{\alpha_k})$. The choice of 20 random rotations and of the 80th percentile were reached after numerical testing, which showed that these choices provided a reasonably accurate estimate of the self-distance of the texture whose realization is \mathbf{x} . In a nutshell, $\text{diam}_\mathbf{x}(j)$ is an efficient Monte Carlo estimate of the baseline computational errors that will affect the computation of $\text{Rdist}_j(\mathbf{x}, \mathbf{y})$. We remark that the Haar measure of $SO(3)$ has to be taken into account when the 20 rotations α_k are randomly selected.

C. Discrimination Between 3-D Textures

The rotationally invariant 3-D texture distance Rdist_j measures the stochastic dissimilarity of two textures. For practical

purposes, we must use Rdist_j at one or more scales j to discriminate two 3-D textures \mathbf{X} and \mathbf{Y} . This question is equivalent to probing whether \mathbf{Y} is statistically identical with a 3-D rigid motion of \mathbf{X} . Recall that even the computable Rdist_j is by definition shift invariant with respect to translates by $\mathbf{t} \in 2^{-j}\mathbb{Z}^3$ because it is extracted by estimating the autocovariance function of \mathbf{x} , which is not affected by the respective translates of the texture. Obviously, one can easily propose more than one texture discrimination rule based on Rdist_j . To this end, we define the discriminant function RD_j by

$$RD_j(\mathbf{x}, \mathbf{y}) = \max \left\{ 0, \frac{\text{Rdist}_j(\mathbf{x}, \mathbf{y}) - \text{diam}_\mathbf{y}(j)}{\text{diam}_\mathbf{y}(j)} \right\}. \quad (23)$$

First we remark that $RD_j(\mathbf{x}, \mathbf{y})$ may not be equal to $RD_j(\mathbf{y}, \mathbf{x})$ due to various numerical errors in the computation of Rdist_j and because $\text{diam}_\mathbf{x}(j)$ may not be equal to $\text{diam}_\mathbf{y}(j)$. Intuitively, the value of the discriminant function $RD_j(\mathbf{x}, \mathbf{y})$ essentially compares how many times further away the most proximal and relatively best 3-D aligned texture realization of \mathbf{y} appears to be from every 3-D rigid-motion-induced realization of \mathbf{x} , where the “unit of the distance measurement” is the self-distance of \mathbf{y} . Since the value of this and of every discriminant function based on Rdist_j is not practically affected by the 3-D rigid motions of each pair of texture realizations that we compare, we can utilize RD_j to devise 3-D rigid-motion-invariant texture discrimination rules. Nevertheless, these rules must be application specific. With this in mind, we coin two rules indicative of the various possibilities.

Rule 1: \mathbf{x} and \mathbf{y} are stochastically the same 3-D texture modulo rigid motions at scale j if $RD_j(\mathbf{x}, \mathbf{y}) = RD_j(\mathbf{y}, \mathbf{x}) = 0$ and statistically different if $RD_j(\mathbf{x}, \mathbf{y}) > 0$ or $RD_j(\mathbf{y}, \mathbf{x}) > 0$.

More conservative variants of rule 1 utilize a “user-defined” threshold $\epsilon > 0$ to declare that \mathbf{x} and \mathbf{y} are realizations modulo 3-D rigid motions of stochastically different 3-D textures at scale j if $RD_j(\mathbf{x}, \mathbf{y}) > \epsilon$ or $RD_j(\mathbf{y}, \mathbf{x}) > \epsilon$. An even more conservative rule would declare the textures as different if both $RD_j(\mathbf{x}, \mathbf{y})$ and $RD_j(\mathbf{y}, \mathbf{x}) > \epsilon$, and certainly more variants of rule 1 can be proposed. Rule 1 can be applied if we have only two texture realizations. Finally, if we wish to apply these rules for multiple scales, we can declare that \mathbf{x} and \mathbf{y} are different if at any of the scales $j = 0, -1, -2, -3$, rule 1 or any rule of choice declares them as different. This approach to multiscale texture discrimination is justified by the fact that we fit very simple texture models to textures whose neighborhood order may or is typically higher than one. By examining the decomposition of the texture output at various scales, we capture voxel interactions of a certain order, which, for $j < 0$, is higher than 1 and cannot be captured if we use our GMRF first-order probe at the original scale only ($j = 0$). The multiscale application of the discrimination rules will become more clear in the experiment section.

Now, suppose that we have more than one texture realization, e.g., $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{r_1}$ are all from the same texture type, e.g., X ; and $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{r_2}$ are all from different texture type, e.g., Y . This is a scenario suitable for an application domain, e.g., all these realizations are texture samples corresponding to two distinct types of soft tissues imaged in a 3-D current transformer (CT) scan. In such a case, tissues of the same histological type

may give rise to texture realizations with quite similar stochastic properties but for which one cannot assert with certainty that they are generated by the same stochastic texture model. To this end, consider the ensembles $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{r_1}$ and $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{r_2}$ as a set of prototypes for these two different types of 3-D textures. Let \mathbf{z} be another texture realization obtained from the same 3-D image. The following is a rule for 3-D rigid-motion-invariant binary classification of 3-D textures.

Rule 2: Let \mathbf{z} be another texture realization extracted from the same 3-D image and assume that \mathbf{z} belongs to one of the two texture types X or Y . We say that \mathbf{z} is of texture type X if $\max\{RD_j(\mathbf{x}_i, \mathbf{z}) : 1 \leq i \leq r_1\} < \min\{RD_j(\mathbf{y}_i, \mathbf{z}) : 1 \leq i \leq r_2\}$. Otherwise, we say that \mathbf{z} is of texture type Y .

We remark that RD_j is not the only choice of a discriminant function. In fact, if one replaces the ratio in the right-hand side of (23) with

$$\frac{Rdist_j(\mathbf{x}, \mathbf{y}) + Rdist_j(\mathbf{y}, \mathbf{x}) - 2 \max\{\text{diam}_x(j), \text{diam}_y(j)\}}{2 \max\{\text{diam}_x(j), \text{diam}_y(j)\}}$$

then the discriminant function becomes symmetric, but this choice no longer works well for texture classification defined by Rule 2. Rule 2 compares a test realization against prototypes, and the Monte Carlo estimate of the numerical errors in the computation of $Rdist_j(\mathbf{x}, \mathbf{y})$ is $\text{diam}_y(j)$, rendering the symmetry of RD_j meaningless unless there is only one prototype from each class and one probe.

VIII. GENERATION OF SYNTHETIC TEXTURES

To test our 3-D rigid-motion-invariant multiscale discrimination approach, we have generated two types of synthetic 3-D textures. The first small set of 3-D textures involves only stochastic textures generated by stationary GMRF models to simply provide a consistency check for our discrimination algorithms within the universe of stationary random textures. We use the algorithm detailed in [41, Ch.4] for generating these textures.

For our second set of numerical tests, we synthesize 3-D textures using cooccurrence matrices [24]. We needed to generate this large test set of arbitrary synthetic 3-D textures at reasonable computational cost. We decided to use 3-D cooccurrence matrix models since these synthetic texture models have been extensively used in computer graphics to emulate extended ranges of real-life 2-D textures extracted from arbitrary photographic data and, in general, are not GMRFs. We deliberately decided not to use GMRF stochastic models for the large test set to demonstrate the classification ability of our algorithm on more general textures, although, we only use GMRF as the probing tool.

Next, we briefly describe our algorithm for generating these textures. In the cooccurrence matrix characterization of a 3-D texture \mathbf{x} , indexed by a finite lattice $\Lambda \subset \mathbb{Z}^3$, the gray levels of \mathbf{x} are usually binned together to create fewer levels for computational efficiency. Call g as the number of gray levels of the image \mathbf{x} after this preliminary compression. Select any fixed symmetric finite subset OFF of \mathbb{Z}^3 , which will be called a set of “offsets.” To each offset $\mathbf{l} \in \text{OFF}$, one typically associates

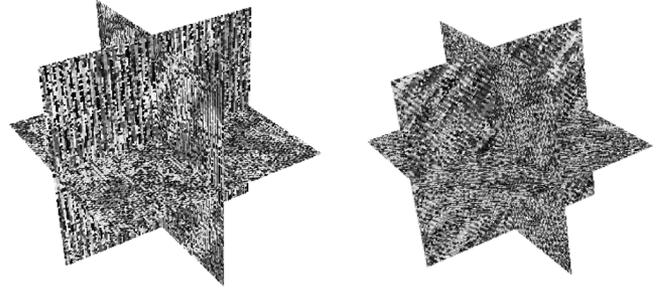


Fig. 2. Examples of 3-D textures obtained from cooccurrence-matrix-based texture synthesis.

the $g \times g$ cooccurrence matrix $C(\mathbf{l})$ with entries $C(\mathbf{l})_{i,j}$ defined by

$$C(\mathbf{l})_{i,j} = \text{cardinal}\{(\mathbf{k}, \mathbf{m}) \in (\Lambda \times \Lambda) \text{ such that } |\mathbf{k} - \mathbf{m}| = |\mathbf{l}|; \mathbf{x}_{\mathbf{k}} = i; \mathbf{x}_{\mathbf{m}} = j\}. \quad (24)$$

We fix a set of $g = 8$ gray levels. To generate each test volume of a 3-D texture, we first select an arbitrary set $\text{OFF} \subset \mathbb{Z}^3$ of five offsets $\mathbf{l} \in [-3, 3]^3$. Then, for each $\mathbf{l} \in \text{OFF}$ we select an arbitrary cooccurrence matrix $C_{\mathbf{l}}$ of size 8×8 , randomly pick eight of $C_{\mathbf{l}}$'s entries, and set them approximately 100 times larger than the rest of the entries of $C_{\mathbf{l}}$. To synthesize a 3-D texture realization \mathbf{x} having those five arbitrary cooccurrence matrices $C(\mathbf{l})$, we extend the texture synthesis algorithm outlined for 2-D-textures in [33] to 3-D. This modified algorithm is a variant of the Metropolis algorithm adapted to minimize the cooccurrence disparity cost function as follows:

$$\text{Cost}(J) := \sum_{\mathbf{l} \in \text{OFF}} \|C_{\mathbf{l}} - K_{\mathbf{l}}\|$$

where J is an arbitrary image and $K_{\mathbf{l}}$ is the cooccurrence matrix of the image J for the offset \mathbf{l} , and the norm of a matrix is the sum of the absolute values of all of its entries.

In contrast to the stochastic descent algorithm in [33], at each iteration, we only update J if the update decreases $\text{Cost}(J)$. This deterministic cost descent differs slightly from the classical stochastic Metropolis algorithm in [33] but converges much faster to a very good local minimum of $\text{Cost}(J)$ quite adequate for the synthesis of arbitrary 3-D textures with eight gray levels and five preassigned cooccurrence matrices. Finally, the dynamic gray-level range of each 3-D image J generated with this algorithm is extended from 8 to 128 gray levels to give J a more “natural” look. This is done by replacing each voxel value $i = 0, 1, \dots, 7$ by a random value in $16 \times i$ and $16 \times (i + 1) - 1$. Realizations of such 3-D textures are shown in Fig. 2.

IX. EXPERIMENTAL RESULTS

To study the performance of our method, we performed three sets of experiments using synthetic and X-ray CT data sets. More details on these experiments and additional experiments are presented in [47].

A. Synthetic Textures Generated by GMRFs

Our first set of experiments intends to perform a first simple consistency check for the proposed rotationally invariant dis-

TABLE I
 $RD_0(\mathbf{x}_i, \mathbf{x}_{i'})$

	\mathbf{X}_1	\mathbf{X}_2	\mathbf{X}_3	\mathbf{X}_4	\mathbf{X}_5	\mathbf{X}_6	\mathbf{X}_7
\mathbf{X}_1	0	193.0	12.6	160.3	161.7	279.0	187.3
\mathbf{X}_2	191.7	0	8.0	895.0	896.0	1455.0	975.3
\mathbf{X}_3	122.0	82.7	0	594.7	597.7	947.5	633.7
\mathbf{X}_4	162.0	893.3	63.8	0	0	2.5	1.3
\mathbf{X}_5	163.0	899.0	64.0	0	0	3.0	1.3
\mathbf{X}_6	186.6	968.0	68.4	1.3	1.3	0	0
\mathbf{X}_7	190.0	980.3	69.5	1.7	1.7	0	0

crimination between 3-D textures. We only use rotations that map the coordinate axes onto themselves because such rotations leave the texture properties invariant. We generate realizations of four synthetic textures \mathbf{X}_4 , \mathbf{X}_5 , \mathbf{X}_6 , and \mathbf{X}_7 using first-order GMRFs with the following parameter vectors θ :

$$\begin{aligned} \mathbf{X}_4 &: (0.1, 0.1, 0.25) & \mathbf{X}_5 &: (0.1, 0.25, 0.1) \\ \mathbf{X}_6 &: (0.05, 0.15, 0.20) & \mathbf{X}_7 &: (0.20, 0.15, 0.05). \end{aligned}$$

Thus, \mathbf{X}_5 is the same as \mathbf{X}_4 rotated by $\pi/2$ around the X -axis, and \mathbf{X}_7 is the same as \mathbf{X}_6 rotated by $\pi/2$ around the Y -axis. We also generate three other realizations of textures, i.e., \mathbf{X}_1 , \mathbf{X}_2 , \mathbf{X}_3 , by using higher-order GMRFs. Table I lists the values of the 3-D rigid-motion-invariant discriminant $RD_0(\mathbf{x}_i, \mathbf{x}_{i'})$ between arbitrary pairs $(\mathbf{x}_i, \mathbf{x}_{i'})$ of these seven 3-D textures at scale $j = 0$. Notice that, in Table I and in all other tables, i denotes row and i' denotes column.

We observed in this experiment that the “distance” $Rdist_j$ is not strictly symmetric but is still nearly very symmetric. This is because our numerical scheme for calculating 3-D texture distances relies on an approximate minimization and a discretization of $SO(3)$ with a fairly moderate density. Samples of rotated autocovariances are approximated by resampling the desired rotation of the autocovariance function on the grid $\mathbb{Z}^3/2$ using the reconstruction part of the fast isotropic wavelet algorithm, as mentioned in Section V-A. We also observed that this resampling technique on the rotated grid $\mathbb{Z}^3/2$ gives better texture discrimination results than linear interpolation. We remark that the lattice $\mathbb{Z}^3/2$ is good for this example, but for other more demanding textures, one may need to use denser lattices such as $\mathbb{Z}^3/2^j$ with $j = 2, 3$. Using the strictest of the variations of rule 1 applied with $\epsilon = 1$ (see Table I) \mathbf{X}_4 , \mathbf{X}_5 and \mathbf{X}_6 , \mathbf{X}_7 are found to be pairs of identical textures up to 3-D rigid motions. Moreover, group \mathbf{X}_4 , \mathbf{X}_5 is not the same (modulo 3-D rigid motions) as the group \mathbf{X}_6 , \mathbf{X}_7 . However, the “distance” RD_0 between the pair \mathbf{X}_4 , \mathbf{X}_5 and the pair \mathbf{X}_6 , \mathbf{X}_7 is not of the same order as the “distance” between either of the \mathbf{X}_4 , \mathbf{X}_5 ; \mathbf{X}_6 , \mathbf{X}_7 ; and \mathbf{X}_1 , \mathbf{X}_2 , and \mathbf{X}_3 . This disparity between the orders of magnitude of these “distances” reflects the structural difference between these textures: The latter three of them are realizations of GMRFs of order higher than 1, whereas the former four are all realizations of first-order GMRFs; therefore, statistically, the disparity among the former four textures should not be as pronounced as the disparity between any of the \mathbf{X}_4 , \mathbf{X}_5 ; \mathbf{X}_6 , \mathbf{X}_7 ; and any of \mathbf{X}_1 , \mathbf{X}_2 , and \mathbf{X}_3 . This observation seems not to be accidental. As the results in the following sections reveal, the proposed “distance” shows discriminative power even

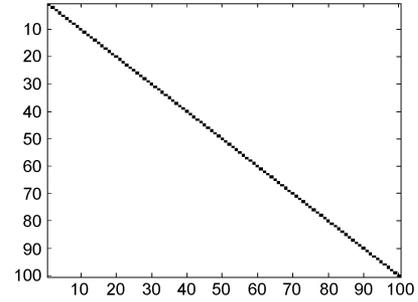


Fig. 3. Discrimination results from combining all three scales using rule 1' for the 4950 pairs of 3-D textures. Texture realization pairs are classified as identical (black pixels) or distinct (white pixels).

when the 3-D texture realizations come from more complex statistical models than GMRFs of the first order. Moreover, our “texture distance probe” helps identifying the five different texture classes in which these seven realizations belong to. Specifically, each of \mathbf{X}_1 , \mathbf{X}_2 , and \mathbf{X}_3 forms a class of its own, and \mathbf{X}_4 , \mathbf{X}_5 and \mathbf{X}_6 , \mathbf{X}_7 form the remaining two classes.

B. Synthetic 3-D Textures Generated by Cooccurrence Matrices

We performed various tests to validate that our approach achieves a high accuracy rate when it discriminates between 3-D textures. With those experiments, we validate that our 3-D texture discrimination is invariant to rigid motions. In the first experiment, we use a set of 100 distinct synthetic 3-D texture realizations generated by distinct arbitrary cooccurrence matrices, as described in Section VIII. We apply to each one of these 3-D texture realizations $\mathbf{x} = \mathbf{x}^{(0)}$ the IMRA decomposition algorithm, to derive the two coarser scale realizations $\mathbf{x}^{(-1)}$ and $\mathbf{x}^{(-2)}$ (see Section VII-A). To demonstrate the discrimination capabilities of the combination of 3-D rigid-motion-invariant texture “distances” $Rdist_j$, monoscale results are combined. For discrimination, we use rule 1', which is a variant of rule 1, as follows: Texture realizations \mathbf{x}_i and $\mathbf{x}_{i'}$ are declared different at scale j modulo 3-D rigid motions if $RD_j(\mathbf{x}_i, \mathbf{x}_{i'}) > 0$ and $RD_j(\mathbf{x}_{i'}, \mathbf{x}_i) > 0$, for $1 \leq i, i' \leq 100$. Using this rule, we say that two textures are different if they are different for any scale j ; otherwise, they are classified as being of the same type. If the pixel (i, i') in Fig. 3 is black, then the realization \mathbf{x}_i is of the same texture type as that of a 3-D rigid motion of $\mathbf{x}_{i'}$. Remarkably, the discrimination accuracy reaches 100% when discrimination results from all three scales are combined, despite the fact that monoscale discrimination accuracy is always strictly less than 100%.

The computation of $Rdist_j$ for each pair requires about 3 min. Computing $Rdist_j$ for 10 000 pairs of 3-D textures at scales $j = 0, -1, -2$ and performing the entire experiment on a 50-processor parallel computing cluster takes less than two days.

The next experiment establishes that the proposed method of 3-D texture discrimination is truly rotationally invariant. To this end, we generate a set of ten synthetic texture realizations of size 240^3 using cooccurrence matrices, as before (see Fig. 4). Then, for any given texture realization, we randomly select a set of eight Euler angles in 3-D by taking two random angles, each for $\alpha \in [10^\circ, 170^\circ]$, $\beta \in [10^\circ, 80^\circ]$, and $\gamma \in [10^\circ, 170^\circ]$. The

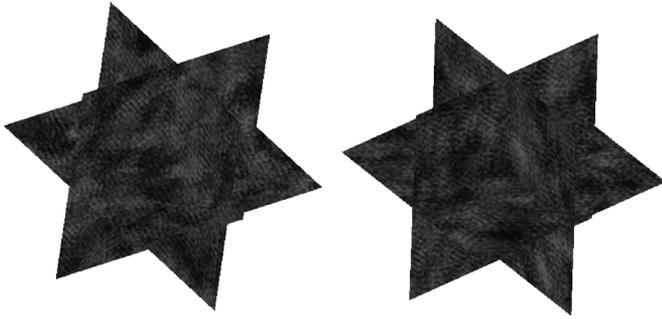


Fig. 4. Example of a texture realization and of its 3-D rotation from one of the ten groups used in the second experiment of Section IX-B.

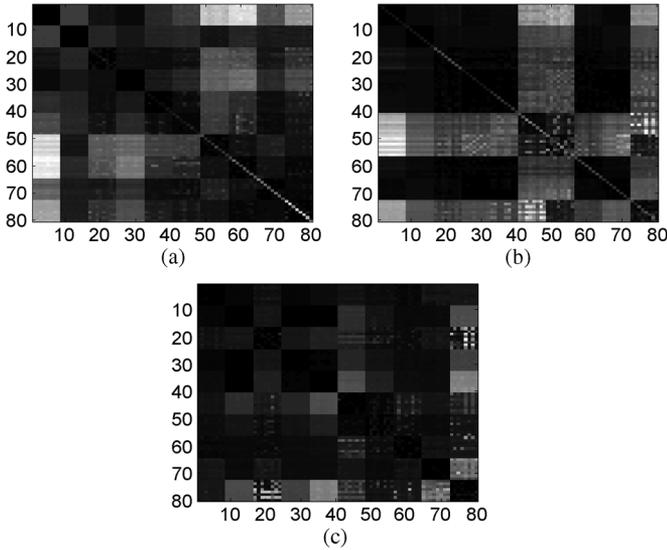


Fig. 5. Grayscale plots of the rotationally invariant distance $Rdist_j(\mathbf{x}_i, \mathbf{x}_{i'})$ when $i \neq i'$ for scales $j = 0, -1, -2$. We use ten 3-D texture realizations, with eight rotations of each realization. The diagonal entries of these plots give the self-distances $diam_{\mathbf{x}_i}(j)$ of these 80 3-D texture realizations. Black pixels encode zero distances, whereas brighter ones encode higher distance values. (a) $Rdist_0$. (b) $Rdist_{-1}$. (c) $Rdist_{-2}$.

range $[10^\circ, 170^\circ]$ is selected to avoid creating almost aligned realizations. Then, we perform one iteration of the IMRA reconstruction algorithm (see Section V-A). The filtered output has size 480^3 . The rotation is implemented by rotating the original 240^3 sampling grid by the appropriate Euler angle and by resampling the filtered 480^3 -voxel output at the points of the rotated 240^3 grid. The use of the IMRA filter h_0 guarantees that numerical errors resulting from the resampling process are well controlled. The value of each voxel of the rotated 240^3 grid is computed from the values of its nearest neighbors in the 480^3 grid using trilinear interpolation. Recall that a similar method is used to compute the rotations of autocovariance matrices during the extraction of 3-D texture signatures. Finally, the rotated texture realization is a cube of size 120^3 extracted from the center of the cube of size 240^3 that we produced in the previous step. This generates a gallery of 80 texture realizations organized in ten groups of eight texture realizations each. All eight members of each group are 3-D rotations of one another.

Applying $Rdist_j$ for $j = 0, -1, -2$ gives the results shown in Fig. 5. Note that images in Fig. 5(a)–(c) are split in 8×8 blocks. Within each block, all pixels have comparable grayscale values.

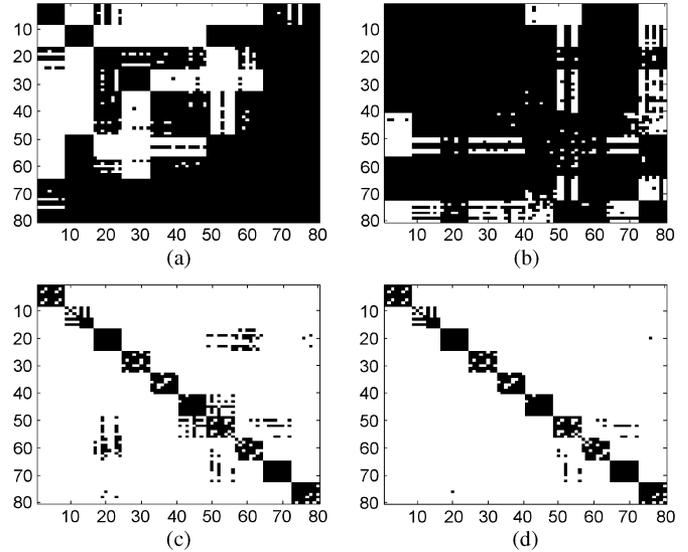


Fig. 6. Discrimination results using rule $1'$ for the ten groups of eight rotated texture realizations. Texture realization pairs are classified as identical (black pixels) or distinct (white pixels). The discrimination accuracy at each scale is given in parenthesis. (a) Scale 0 (accuracy of 65.13%). (b) Scale -1 (accuracy of 53.54%). (c) Scale -2 (accuracy of 97.41%). (d) Combine all three scales (accuracy of 98.75%).

The pixels belonging to the same 8×8 block in the main diagonal of each plot appear to have the same grayscale value indicating that the self-distances of the eight rotated realizations of the same texture are almost equal. This observation underscores the expected result: Since $Rdist_j$ is approximately rotationally invariant, $diam_{\mathbf{x}}(j)$ is almost equal to $diam_{\mathbf{x}\alpha}(j)$ ($\alpha \in SO(3)$). Applying now the discrimination rule $1'$ per scale and by combining scales as we did in Experiment 1 gives the results shown in Fig. 6. Inspection of the latter plot shows that the use of rule $1'$ enables a very good rotation invariant discrimination, particularly effective in deciding that two textures are not 3-D rigid-motion equivalent. If, instead of rule $1'$, we use rule 1, then we observe roughly the same high rate of discrimination accuracy, but the number of wrong hits increases in the off-diagonal blocks and decreases in the string of the 8×8 diagonal blocks.

We perform a similar experiment to show that the proposed texture discrimination method is 3-D translation invariant as well. Again, we generate a set of ten synthetic texture realizations as in the previous experiment. Each of them has size 240^3 . Eight cubes of size 120^3 are randomly extracted from each octant of all ten volumes. This provides us with eight translations for each texture realization. By applying rule $1'$ in three scales $j = 0, -1, -2$ and by combining discrimination results shown in Fig. 7, the discrimination accuracy reaches 99.76%. The combined application of discrimination rule $1'$, as given previously, yields results plotted in Fig. 7.

C. Shot of Soft-Tissue Discrimination

In this section, we show how the proposed method is capable of soft-tissue discrimination in 3-D scans when tissues are associated with 3-D textures. We use two different series of abdomen X-ray CT scans acquired from the same patient at different times during the same study. Gall bladder and muscle are

TABLE II

SERIES I: VALUES OF $RD_{-1}(\mathbf{x}, \mathbf{y})$. G_1, \dots, G_4 AND M_1, \dots, M_4 ARE USED AS PROTOTYPES FOR GALL BLADDER AND MUSCLE, RESPECTIVELY. TG AND TM ARE TEST SAMPLES FOR GALL BLADDER AND MUSCLE. OBSERVE THAT THE VALUES IN COLUMN OF TG AND TM INDICATES THAT TG IS MORE “CLOSE” TO G_1, \dots, G_4 , AS COMPARED WITH M_1, \dots, M_4 . SIMILARLY, TM IS MORE “CLOSE” TO M_1, \dots, M_4 AS COMPARED WITH G_1, \dots, G_4

	G_1	G_2	G_3	G_4	TG	M_1	M_2	M_3	M_4	TM
G_1	0	18.1	4.5	0	4.2	697.4	14.9	1811.1	14.3	162.3
G_2	23.2	0	32.5	0.7	15.7	1552.9	19.4	3381.8	22.8	256.9
G_3	1.5	33.4	0	0	2.6	875.2	17.2	2134.2	17.0	195.9
G_4	74.1	71.1	21.4	0	10.3	1995.7	25.2	4389.8	30.4	309.2
TG	50.1	53.5	7.3	0	0	1658.1	23.9	3522.4	24.5	283.9
M_1	116.9	93.9	55.7	1.8	30	0	1.9	105.4	4.0	27.4
M_2	205.2	126.3	90.6	2.7	41.5	569.6	0	572.4	1.5	10.8
M_3	200.7	115.7	84.1	2.4	40.7	56.8	0.5	0	2.1	15.9
M_4	217.4	107.6	92.6	2.3	46.2	701.6	0	666.1	0	48.7
TM	250.9	128.5	101.3	2.8	49.0	488.2	0.1	293.3	3.8	0

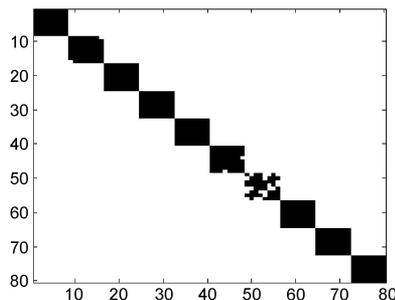


Fig. 7. Combined discrimination results using rule 1' for the ten groups of eight spatially shifted texture realizations. Texture realization pairs are classified as identical (black pixels) or distinct (white pixels). The discrimination accuracy reaches 99.76%.

the two tissue types we consider in each of the two CT-scan series. We show that we can distinguish these two types of soft tissue by using the 3-D texture discrimination techniques we develop. Suppose that \mathbf{x} and \mathbf{y} are texture realizations from gall bladder and/or muscle tissue.

First, consider the standardized mean difference $D_M(\mathbf{x}, \mathbf{y})$ defined by

$$D_M(\mathbf{x}, \mathbf{y}) = \frac{|\text{Mean}(\mathbf{x}) - \text{Mean}(\mathbf{y})|}{\sqrt{\text{Var}(\mathbf{x}) + \text{Var}(\mathbf{y})}}$$

where $\text{Var}(\mathbf{x})$ and $\text{Var}(\mathbf{y})$ are sample variances of \mathbf{x} and \mathbf{y} , respectively. This expression gives a simple measure of how much the average intensity values of \mathbf{x} in Hounsfield units are different from the average intensities of \mathbf{y} in a standardized “unit of measurement,” which can be considered as the Z-score of the means of \mathbf{x} and \mathbf{y} , assuming that these two tissue samples are independent of one another. From each CT-scan series, we extract five texture realizations from each tissue type, and we designate four of them, all from the same tissue type as a training set (referred to as G_i and M_i , $i = 1, 2, 3, 4$) and the remaining one of them as a test sample (referred to as TM and TG ; “ G ” signifies gall bladder). Each of these sample volumes has a size of 23^3 . We observe that, in both CT-scan series, $D_M(x, y)$ drops below 2, which implies that we cannot safely distinguish these two tissue types by using only the standardized difference of average intensities or by thresholding Hounsfield unit intensities in general.

Tables II and III show the values of RD_{-1} for different pairs of tissue samples from both CT-scan series. First, observe that RD_{-1} takes positive values even for 3-D texture pairs of the

same tissue type. By no means this should be interpreted as difference in tissue types. Natural and anatomical variations among different samples of the same tissue may void the spatial homogeneity assumption for textures forcing RD_{-1} to assume positive values. For instance, a volume in the gall bladder close to a cystic artery has an extra directional feature due its proximity to the artery that is full of contrast, which other regions of the gall bladder, less proximal to this artery, do not have. This observation shows that rule 2 in Section VII-C and its variants facilitate the correct tissue discrimination. Rule 2 introduces the use of training tissue samples. Observe that the values in columns TG and TM in both tables indicate that TG is more “close” to G_1, \dots, G_4 than to M_1, \dots, M_4 and that the opposite holds true for TM . The test samples from both tissue types are, indeed, correctly classified as “gall bladder” or “muscle” using rule 2 applied on the values of RD_{-1} shown in Tables II and III.

In conclusion, the statistical disparity RD_j applied to a single scale or to a combination of scales can be used for binary classification purposes as in this case of soft-tissue classification, when intensity value thresholding is not statistically sufficient for tissue classification. In fact, if both are combined, the classification accuracy should improve.

X. CONTRIBUTIONS, SUMMARY AND FUTURE WORK

We present an innovative approach to rotationally invariant classification of 3-D textures. Instead of adopting an isotropic model for the texture statistics, we define a 3-D rigid-motion-invariant distance for high-dimensional GMRF-based texture signatures, which are rotationally covariant and invariant to 3-D translations. Therefore, the proposed 3-D texture signatures are indexed by the full group of 3-D rotations. In practice, they are built by fitting simple GMRF models to each rotation of a given texture, despite that those textures may not be GMRFs. The GMRF models associated to each rotation of a 3-D texture are simply used as a probing tool, to encode statistics of local interactions between neighboring pixels into low-dimensional feature vectors computable at low cost. Our main contributions are the following:

- 1) A novel rotationally invariant distance between the high-dimensional signatures of two arbitrary 3-D textures \mathbf{x} and \mathbf{y} ; this distance is rigorously based on calculating the Haar integral (with respect to the measure of $SO(3)$) of the KL

TABLE III

SERIES 2: VALUES OF $RD_{-1}(\mathbf{x}, \mathbf{y})$, G_1, \dots, G_4 AND M_1, \dots, M_4 ARE USED AS A TRAINING SET FOR GALL BLADDER AND MUSCLE, RESPECTIVELY. TG AND TM ARE TEST SAMPLES FOR GALL BLADDER AND MUSCLE. OBSERVE THAT, ONCE AGAIN, THE VALUES IN COLUMN OF TG AND TM INDICATES THAT TG IS MORE "CLOSE" TO G_1, \dots, G_4 AS COMPARED WITH M_1, \dots, M_4 SIMILARLY, TM IS MORE "CLOSE" TO M_1, \dots, M_4 , AS COMPARED WITH G_1, \dots, G_4

	G_1	G_2	G_3	G_4	TG	M_1	M_2	M_3	M_4	TM
G_1	0	0	2.3	5.6	5.0	162.2	533.8	235.3	1504.2	59.7
G_2	115.0	0	0	20.2	0.6	329.8	794.7	304.2	2297.1	82.3
G_3	18.7	0	0	8.9	0.2	181.1	495.0	199.6	1435.3	59.1
G_4	59.7	0.2	5.8	0	6.4	28.7	284.0	124.4	535.9	34.4
TG	82.1	0	1.5	17.1	0	368.2	791.4	312.5	2394.3	80.9
M_1	91.4	0.4	16.5	1.1	12.3	0	126.3	35.2	130.4	15.8
M_2	227.0	2.0	35.0	17.8	22.6	132.6	0	0.3	104.0	6.4
M_3	228.5	1.9	36.7	16.8	21.4	89.3	18.4	0	83.0	10.1
M_4	185.6	1.4	30	9.5	17.7	29.7	21.2	6.8	0	8.9
TM	234.9	1.9	39.4	19.0	20.2	156.4	111.8	15.0	252.5	0

divergences between pairs of GMRF models associated to each 3-D rotation of textures \mathbf{x} and \mathbf{y} .

- 2) Implementation of 3-D rotations of a 3-D texture in a computationally efficient manner. The idea of rotating the autocovariance function corresponding to the texture, instead of the texture itself, makes the 3-D rotations computationally tractable. This novelty enables discrimination between 3-D textures having a broad variety of directional characteristics.
- 3) An efficient method for the computation of the rotationally invariant distance by means of the Euler angle parameterization of $SO(3)$ and the selection of a moderate-size finite grid of points in $SO(3)$ uniformly distributed with respect to the Haar measure of $SO(3)$. The block-circulant structure of the covariance matrix of the GMRF model enhances the computational efficiency since it allows the calculation of the KL distance in the Fourier domain.
- 4) Rules for the 3-D rigid-motion-invariant texture discrimination and binary classification, which take into account the proposed 3-D rigid-motion-invariant distance between textures computed at a range of scales. These rules allow the use of simple GMRF models and avoid computationally costly parameter estimation of GMRFs with long-range interactions. We experimentally establish that these rules enhance the sensitivity of rigid-motion-invariant discrimination between 3-D textures and that they are applicable even to non-GMRF textures.

Our motivation to study stochastic 3-D textures comes from medical image analysis. X-ray CT and MRI scanners increasingly acquire high-spatial-accuracy 3-D images in which textures are generated by the response of soft biological tissues to the imaging modality. Those 3-D textures are not suitably modeled by deterministic approaches for two reasons. First, natural local tissue variations modify 3-D textures corresponding to the same type of tissue, which is a fact that is quite well shown with the variations of sample mean values among tissue samples of the same type (see Section IX-C); second, image acquisition "noise" contributes to the variability of these textures.

Although, we do not claim that the herein proposed methods for 3-D texture discrimination are applicable for all types of soft-tissue discrimination/classification in CT and MRI 3-D images, we provide in Section IX-C experimental evidence establishing a method that can be utilized for at least binary soft-

tissue classification in medical 3-D image analysis. In X-ray CT in particular, blood-infused tissues have very similar scattering properties due to their high content of water molecules. As a result, several types of soft tissues are indiscriminable by intensity thresholding only. The last experiment in the previous section shows that soft-tissue discrimination in the native dimensionality of X-ray CT images is a valid problem. In this context, 3-D rigid-motion insensitive tissue discrimination is very desirable as soft tissues occupy the 3-D space in a random variety of orientations and are often associated with nonisotropic textures.

A seemingly practical constraint of the proposed method is the size of the 3-D texture patch required for the calculation of the 3-D invariant distances between textures. We have empirically established that we need a 3-D texture patch with a size of $20 \times 20 \times 20$ – $30 \times 30 \times 30$ for accurate 3-D texture discrimination. The need to downsample the image data and to perform accurate multiscale texture discrimination makes this patch-size requirement rather restrictive for tissue types whose spatial extent varies significantly. Alternatively, downsampling can be omitted when discrimination is applied to scales $j < 0$, but then, the first-order neighbors become the pixels two or four, $\dots, 2^j$ apart from the center of the neighborhood. We are currently developing a method [47] on how to modify the feature extraction we have presented in this paper to overcome the need to use rectangular 3-D texture patches.

ACKNOWLEDGMENT

The authors would like to thank the reviewers for the valuable comments that helped us to improve the presentation of this paper. They also thank Dr. G. Gladish of the MD Anderson Cancer Center for providing them with the two anonymized data sets used in the last experiment.

REFERENCES

- [1] "Euler angles," [Online]. Available: http://en.wikipedia.org/wiki/Euler_angles
- [2] S. K. Alexander, R. Azencott, and M. Papadakis, "Isotropic multiresolution analysis for 3D-textures and applications in cardiovascular imaging," in *Wavelets XII*, D. Van Der Ville, V. Goyal, and M. Papadakis, Eds. Bellingham, WA: SPIE, 2007, pp. 67 011S–1–67 011S–12.
- [3] C. Audet and J. E. Dennis, Jr., "Analysis of generalized pattern searches," *SIAM J. Optim.*, vol. 13, no. 3, pp. 889–903, 2002.
- [4] R. Azencott and D. Dacunha-Castelle, *Series of Irregular Observations*. New York: Springer-Verlag, 1986.

- [5] R. Azencott, C. Graffigne, and C. Labourdette, *Edge Detection and Segmentation of Textured Images Using Markov Fields*. New York: Springer-Verlag, 1992, pp. 75–88.
- [6] R. Azencott, J. P. Wang, and L. Younes, “Texture classification using Gabor filters,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 19, no. 2, pp. 148–153, Feb. 1997.
- [7] B. G. Bodmann, M. Papadakis, D. J. Kouri, S. D. Gertz, P. Cherukuri, D. Vela, G. W. Gladish, D. D. Cody, I. Abodashy, J. L. Conyers, J. T. Willerson, and S. W. Casscells, “Frame isotropic multiresolution analysis for micro CT scans of coronary arteries,” in *Wavelets XI*, M. Papadakis, A. Laine, and M. Unser, Eds. Bellingham, WA: SPIE, 2005, pp. 59 1410/1–59 1410/12.
- [8] C. De Boor, R. A. DeVore, and A. Ron, “Approximation from shift-invariant subspaces of $L_2(\mathbb{R}^d)$,” *Trans. Amer. Math. Soc.*, vol. 341, no. 2, pp. 787–806, Feb. 1994.
- [9] C. Bouman and B. Liu, “Multiple resolution segmentation of textured images,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-13, no. 2, pp. 99–113, Feb. 1991.
- [10] T. Chang and C. Kuo, “Texture analysis and classification with tree-structured wavelet transform,” *IEEE Trans. Image Process.*, vol. 2, no. 4, pp. 429–441, Oct. 1993.
- [11] R. Chellappa, “Two-dimensional discrete Gaussian Markov random field models for image processing,” in *Progress in Pattern Recognition*. Amsterdam, The Netherlands: Elsevier, 1985, pp. 79–112.
- [12] R. Chin and C. Harlow, “Automated visual inspection: A survey,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-4, no. 6, pp. 557–573, Nov. 1982.
- [13] G. S. Chirikjian and A. B. Kyatkin, *Engineering Applications of Non-commutative Harmonic Analysis: With Emphasis on Rotation and Motion Groups*. Boca Raton, FL: CRC Press, 2000.
- [14] H. Choi and R. Baraniuk, “Multiscale image segmentation using wavelet domain hidden Markov models,” in *Proc. SPIE Conf. Math. Model., Bayesian Estimation, Inverse Problems*, 1999, vol. 3816, pp. 306–320.
- [15] F. S. Cohen, Z. Fan, and M. A. Patel, “Classification of rotated and scaled textured images using Gaussian Markov Random Field models,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 2, pp. 192–202, Feb. 1991.
- [16] G. R. Cross and A. K. Jain, “Markov random field texture models,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-5, no. 1, pp. 25–39, Jan. 1983.
- [17] H. Derin and H. Elliott, “Modeling and segmentation of noisy and textured images using Gibbs random fields,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-9, no. 1, pp. 39–55, Jan. 1987.
- [18] D. Dunn and W. E. Higgins, “Optimal Gabor filters for texture segmentation,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 4, no. 7, pp. 947–964, Jul. 1995.
- [19] I. M. Elfadel and R. W. Picard, “Gibbs random fields, cooccurrences, and texture modeling,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 16, no. 1, pp. 24–37, Jan. 1994.
- [20] K. Etamad and R. Chellappa, “Separability based tree-structured local basis selection for texture classification,” in *Proc. ICIP*, Austin, TX, 1994, vol. 3, pp. 441–445.
- [21] J. Fehr and H. Burkhardt, “3-D rotation invariant local binary patterns,” in *Proc. ICPR*, Tampa, FL, 2008, pp. 1–4.
- [22] S. Geman and D. Geman, “Stochastic relaxation, Gibbs distributions, and the bayesian restoration of images,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-6, no. 6, pp. 721–741, Nov. 1984.
- [23] R. Haralick, “Statistical and structural approaches to texture,” *Proc. IEEE*, vol. 67, no. 5, pp. 786–804, May 1979.
- [24] R. M. Haralick, Dinstein, and K. Shanmugam, “Textural features for image classification,” *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-3, no. 6, pp. 610–621, Nov. 1973.
- [25] S. Jain, M. Papadakis, and E. Dussaud, “Explicit schemes in seismic migration and isotropic multiscale representations,” in *Radon Transforms, Geometry, and Wavelets*, E. Grinberg, D. R. Larson, P. E. T. Jorgensen, P. Massopust, G. Olafsson, E. T. Quinto, and B. Rubin, Eds. Providence, RI: AMS, 2008, pp. 177–200.
- [26] R. L. Kashyap, “Analysis and synthesis of image patterns by spatial interaction models,” in *Progress in Pattern Recognition*. Amsterdam, The Netherlands: Elsevier, 1981, pp. 149–186.
- [27] R. L. Kashyap and R. Chellappa, “Estimation and choice of neighbors in spatial-interaction models of images,” *IEEE Trans. Inf. Theory*, vol. IT-29, no. 1, pp. 60–72, Jan. 1983.
- [28] R. L. Kashyap and A. Khotanazad, “A model-based method for rotation invariant texture classification,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-8, no. 4, pp. 472–481, Jul. 1986.
- [29] V. A. Kovalev, F. Kruggel, H. J. Gertz, and D. Y. von Cramon, “Three-dimensional texture analysis of MRI brain data plural**sets,” *IEEE Trans. Med. Imag.*, vol. 20, no. 5, pp. 424–433, May 2001.
- [30] A. S. Kurani, D. H. Xu, J. D. Furst, and D. S. Raicu, “Co-occurrence matrices for volumetric data,” presented at the 7th IASTED Int. Conf. Comput. Graph. Imag. (CGIM), Kauai, HI, Aug. 16–18, 2004.
- [31] S. Lakshmanan and H. Derin, “Valid parameter space for 2-D Gaussian Markov random fields,” *IEEE Trans. Inf. Theory*, vol. 39, no. 2, pp. 703–709, Mar. 1993.
- [32] T. S. Lee, “Image representation using 2-D Gabor wavelets,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 18, no. 10, pp. 959–971, Oct. 1996.
- [33] G. Lohmann, “Analysis and synthesis of textures: A co-occurrence-based method,” *Comput. Graph.*, vol. 19, no. 1, pp. 29–36, 1995.
- [34] S. Y. Lu and K. S. Fu, “A syntactic approach to texture analysis,” *Comput. Graph. Image Process.*, vol. 7, no. 3, pp. 303–330, Jun. 1978.
- [35] A. Madabhushi, M. Feldman, D. Metaxas, D. Chute, and J. Tomaszewski, “A novel stochastic combination of 3d texture features for automated segmentation of prostatic adenocarcinoma from high resolution MRI,” in *Proc. Med. Image Comput. Comput.-Assist. Interv.*, 2003, vol. 2878, pp. 581–591.
- [36] S. G. Mallat, “Multiresolution approximations and wavelet orthonormal bases of $l^2(\mathbb{R})$,” *Trans. Amer. Math. Soc.*, vol. 315, no. 1, pp. 69–87, Sep. 1989.
- [37] J. Mao and A. K. Jain, “Texture classification and segmentation using multiresolution simultaneous autoregressive models,” *Pattern Recognit.*, vol. 25, no. 2, pp. 173–188, Feb. 1992.
- [38] M. Do and M. Vetterli, “Rotation invariant texture characterization and retrieval using steerable wavelet domain hidden Markov models,” *IEEE Trans. Multimedia*, vol. 4, no. 4, pp. 517–527, Dec. 2002.
- [39] M. Papadakis, B. G. Bodmann, S. K. Alexander, D. Vela, S. Baid, A. A. Gittens, D. J. Kouri, S. D. Gertz, S. Jain, J. R. Romero, X. Li, P. Cherukuri, D. D. Cody, G. W. Gladish, I. Aboshady, J. L. Conyers, and S. W. Casscells, “Texture-based tissue characterization for high-resolution CT-scans of coronary arteries,” *Commun. Numer. Methods Eng.*, vol. 25, no. 6, pp. 597–613, 2009.
- [40] R. Porter and N. Canagarajah, “Robust rotation-invariant texture classification: Wavelet, Gabor filter and GMRF based schemes,” *Proc. Inst. Elect. Eng.—Vis. Image Signal Process.*, vol. 144, no. 3, pp. 180–188, Jun. 1997.
- [41] E. B. Rangelova, “Segmentation of textured images on three-dimensional lattices,” Ph.D. dissertation, Univ. Dublin, Trinity College, Dublin, Ireland, 2002.
- [42] C. C. Reyes-Aldasoro and A. Bhalerao, “Volumetric feature selection for MRI,” in *Proc. Inf. Process. Med. Imag.*, 2003, vol. 2732, pp. 282–293.
- [43] J. R. Romero, S. Alexander, S. Baid, S. Jain, and M. Papadakis, “The geometry and the analytic properties of isotropic multiresolution analysis,” *Adv. Comput. Math.*, vol. 31, no. 1–3, pp. 283–328, Oct. 2009.
- [44] A. Teuner, O. Pichler, and B. J. Hosticka, “Unsupervised texture segmentation of images using tuned Gabor filters,” *IEEE Trans. Image Process.*, vol. 4, no. 6, pp. 863–870, Jun. 1995.
- [45] F. Tomita and S. Tsuji, *Computer Analysis of Visual Textures*. Norwell, MA: Kluwer, 1990.
- [46] M. Unser, “Texture classification and segmentation using wavelet frames,” *IEEE Trans. Image Process.*, vol. 4, no. 11, pp. 1549–1560, Nov. 1995.
- [47] S. Upadhyay, S. Jain, M. Papadakis, and R. Azencott, “3-D-Rigid motion invariant discrimination and classification of 3-D-textures,” in *Proc. SPIE—Wavelets and Sparsity XIV*, M. Papadakis, D. Van De Ville, and V. Goyal, Eds., 2011, vol. 8138, pp. 813 821-1–813 821-17, DOI: 10.1117/12.891721.
- [48] D. H. Xu, A. S. Kurani, J. D. Furst, and D. S. Raicu, “Run-length encoding for volumetric texture,” presented at the 4th IASTED Int. Conf. Vis., Imag. Image Process. (VIIP), Marbella, Spain, Sep. 6–8, 2004.



Saurabh Jain received the B.Sc. degree from St. Stephen's College, Delhi University, Delhi, India, in 2003 and the M.S. and Ph.D. degrees in applied mathematics from the University of Houston, Houston, TX, in 2005 and 2009, respectively.

He currently holds a postdoctoral research fellowship with the Center for Imaging Science, Johns Hopkins University, Baltimore, MD. He has broad industrial and academic research experience with practical problems in medical imaging and seismic imaging. His research interests include wavelet and

harmonic analysis, Markov random fields, and diffeomorphic shape matching and image registration.



Manos Papadakis received the Ph.D. degree in mathematics from the University of Athens, Athens, Greece.

He is currently an Associate Professor with the Department of Mathematics, University of Houston, Houston, TX, and his interests include sparse representations, wavelet analysis in multidimensions, frame theory, and biomedical image analysis. The primary focus of his research is on imaging neuroscience with emphasis on spine detection, classification, and quantification of synaptic plasticity

and homeostasis. Part of his work aims in eliminating errors in computer vision applications resulting from the directional preference of 1-D filter designs and on developing 3-D rigid-motion covariant data representations. He and his collaborators propose that the best antidote for the former problem is the use of isotropic filtering methods, treating all directions in an equitable manner and thus enabling the preprocessing and sampling of multidimensional data sets in their original dimensionality with less artifacts.



Sanat Upadhyay received the B.Sc. degree in applied physical sciences from St. Stephen's College, Delhi University, New Delhi, India, and the M.S. degree in applied mathematics from the University of Houston, Houston, TX, and is currently working toward the Ph.D. degree in mathematics from the University of Houston.

His research interests include 3-D texture classification and multivariable sparse representations with applications to face recognition.



Robert Azencott received the Masters degree from the Ecole Normale Supérieure, Paris, France; in 1964, the MA degree from Harvard University, Cambridge, MA, in 1966; and the Doctorat d'Etat degree in mathematics from the University of Paris 6, Paris, France, in 1970.

He is currently a Professor of mathematics with the University of Houston, Houston, TX, and an Emeritus Professor of mathematics with Ecole Normale Supérieure, Cachan, France. He is an Applied Mathematician. He has directed 50 Ph.D.

students in probability, statistics, automatic learning, multisensors data mining, and image analysis with more than 100 published papers and books. He has held university professorships in the USA (University of California at Berkeley, Brandeis, and Houston), and in Paris, France (Ecole Normale Supérieure, University Paris 7, and Paris 11). With major industrial partners, he has launched applied research projects in aeronautics, nuclear science, satellite imaging, chemistry, automotive, pharmacy, etc. His current research interests include 3-D medical imaging analysis and diffeomorphic matching of deformable shapes.